

# **Decomposing the Capability Overhang: Access, Agency, and the Geography of AI Adoption**

**Ewan Simpson**

**March 27 2026**

## **Abstract**

The gap between what AI systems can do and what people actually do with them, recently termed the "capability overhang" by OpenAI, has become a prominent policy concern at the national and international level. Yet the concept, as currently used, treats this gap as a single undifferentiated phenomenon and proposes correspondingly uniform remedies: more access, more education, more infrastructure. This paper argues that capability overhang is structurally two-dimensional. It simultaneously reflects failures of access: the inability of populations to reach AI tools, and failures of agency: the inability to use those tools in ways that generate productive, iterative engagement. These two failures are distinct in origin, operate at different speeds, and require fundamentally different policy responses.

We introduce the AI Matrix (Simpson, 2025a) as a framework for decomposing capability overhang along these two dimensions, and operationalize it using observed interaction data from the Anthropic Economic Index (AEI), a privacy-preserving analysis of 818,492 Claude.ai Free and Pro consumer conversations across 173 countries during November 13-20, 2025. Using use case composition, collaboration pattern distributions, and task success rates as proxies for the agency dimension, combined with conversation volume as a proxy for effective AI reach within the same dataset, we show that countries cluster into patterns broadly consistent with the AI Matrix's four quadrants. The resulting typology describes three observable stages of AI adoption that represent a progression in structural complexity: an education-first stage characterized by single-domain use concentrated in academic coursework and high rates of passive delegation; a professional adoption stage characterized by single-domain concentration in work tasks among technically skilled populations; and a consumer diffusion stage characterized by broad use across work, personal, and creative purposes with declining rates of passive delegation.

We argue that this typology is not merely descriptive. By mapping it onto the AI Matrix's transition logic, we offer a theoretically grounded interpretation of why the adoption curve takes the shape that Anthropic and OpenAI have independently documented, and why global gaps between countries have remained stable even as within-country convergence accelerates. This interpretation is consistent with the observed evidence rather than established by it. The paper is

explicit about the limitations of a cross-sectional, single-provider dataset: the AEI data captures only countries above the access threshold, cannot establish causal direction, and reflects one week of interactions on one platform. Within these constraints, we argue that the patterns are consistent with the AI Matrix's predictions and provide an empirical foundation for a policy framework that moves beyond access alone.

Keywords: AI adoption, capability overhang, AI Matrix, digital divide, human agency, technology diffusion, global inequality, Anthropic Economic Index

## **1. Introduction**

In January 2026, OpenAI arrived at the World Economic Forum in Davos with a striking argument. The problem with AI, its executives said, was no longer capability. Current models could already outperform most of the tasks that most professionals asked them to do. The problem was use. Between what AI systems could do and what people actually did with them lay a gap so large that closing it had become, in OpenAI's framing, as important as any further advance in the technology itself. They called this gap the capability overhang, and they warned that without deliberate action to close it, the biggest gains from AI would flow to the countries, companies, and workers already testing its limits, while everyone else fell further behind in ways that would be difficult to reverse (OpenAI, 2026a).

The warning landed in fertile ground. Anthropic's parallel program of research, the Anthropic Economic Index, had been documenting the geographic contours of AI adoption across more than 150 countries since January 2025. Its findings were consistent and sobering. Claude usage per capita correlated strongly with GDP per capita. Countries with higher incomes used AI more collaboratively, in ways that were more iterative, more exploratory, and more oriented toward augmenting human judgment. Countries with lower incomes used AI more directly, more narrowly, and more heavily for educational coursework than for professional work or personal use (Appel et al., 2025; Appel et al., 2026a). Microsoft's AI Economy Institute independently documents a widening diffusion divide using telemetry-derived estimates of the share of the working-age population using any generative AI product, adjusted for device and internet penetration (Microsoft AI Economy Institute, 2026). In H2 2025, the Global North reached 24.7 percent adoption against 14.1 percent in the Global South, a gap that grew over the year. Anthropic's characterisation of Microsoft data further specifies that AI use for educational purposes concentrates among lower per-capita income countries while AI use for leisure concentrates among wealthier populations (Appel et al., 2026a). Three lines of evidence converge on the same structural finding, though with different levels of methodological transparency: Anthropic and OpenAI publish primary behavioral data, while Microsoft's AI Economy Institute publishes primary

telemetry-based data on population-level diffusion rates (Microsoft AI Economy Institute, 2026). The measures are complementary rather than equivalent: the AEI data captures how AI is used; the Microsoft data captures how widely it has spread.

What the literature does not provide is a theoretical account of what these patterns mean, why they persist, and what would actually change them. The capability overhang concept, as OpenAI deploys it, identifies a gap without specifying its structure. It describes the distance between AI's potential and its realization as though that distance were a single thing, driven by a single set of causes and responsive to a single family of remedies. The policy prescriptions that follow are not wrong, but they are undifferentiated. They assume that the same intervention closes the gap regardless of where a country sits on the adoption curve, regardless of whether its primary constraint is the availability of AI tools or the capacity to use those tools productively once they are available. That assumption is not supported by the evidence, and it matters enormously for how societies allocate scarce resources.

This paper argues that capability overhang is structurally two-dimensional. It reflects, simultaneously, failures of access and failures of agency. Access refers to the structural availability of AI tools, encompassing the infrastructure, connectivity, affordability, language support, and regulatory conditions that determine whether a given population can reach the technology at all. Agency refers to the capacity to engage with those tools in ways that generate productive outcomes, including the sophistication of interaction, the willingness to iterate rather than simply delegate, and the educational and institutional conditions that support critical rather than passive use. These two dimensions can fail independently. A country can have broad access and low agency, producing widespread but shallow engagement. A country can have concentrated access and high agency among the subset of users who have it, producing intensive professional use that leaves the majority of the population unreached. Each of these configurations describes a different problem, requires a different diagnosis, and calls for a different intervention.

We introduce the AI Matrix (Simpson, 2025a) as a framework for making this decomposition analytically tractable. The AI Matrix organizes countries, institutions, and individuals along two axes, access to AI resources and human agency, producing four quadrants that describe qualitatively distinct relationships with AI technology. These are Full Empowerment, Elite Empowerment, Passive Dependency, and Full Dependency. Crucially, the framework was designed as dynamic rather than static. It specifies not only where actors currently sit but the conditions and mechanisms through which they move between quadrants. This transition logic is what existing accounts of capability overhang and the adoption curve conspicuously lack.

To operationalize the framework empirically, we draw on the Anthropic Economic Index (AEI), a privacy-preserving analysis of 818,492 Claude.ai Free and Pro consumer conversations across 173 countries from November 13-20, 2025, released publicly by Anthropic with an explicit invitation to external researchers (Appel et al., 2026a). The AEI provides, for each country with sufficient observation volume, distributions of use case (work, coursework, personal), collaboration pattern (directive, task iteration, learning, feedback loop, validation), task success rates, and education-level proxies for both prompt and response sophistication. These variables serve as observable proxies for the agency dimension of the AI Matrix. Combined with external indicators of AI access, they allow us to place countries in approximate quadrant positions and to identify the three-stage adoption typology that emerges from the data. We characterize these as an education-first stage consistent with Passive Dependency, a professional adoption stage consistent with Elite Empowerment, and a consumer diffusion stage consistent with movement toward Full Empowerment.

The AEI data is cross-sectional, covers a single week, reflects one platform, and captures only countries already above the access threshold. The countries most in need of the framework are largely invisible in the data, because they are not yet present in it in meaningful numbers. The agency composite we construct from collaboration patterns and task outcomes is a proxy derived from observed interaction behavior, not a validated direct measure of the underlying construct. The paper's empirical contribution is pattern identification and framework validation, not causal inference. We show that the distribution of countries in the data is consistent with the AI Matrix's predictions. We do not claim to have demonstrated the mechanisms through which countries move between quadrants.

That said, the convergence of three lines of evidence on the same descriptive pattern, combined with the consistency of that pattern with a theoretical framework published prior to any of these datasets appearing is of value. The three lines of evidence differ in methodological transparency: Anthropic and OpenAI publish primary behavioral data, while Microsoft's AI Economy Institute measures population-level AI diffusion rates derived from telemetry (Microsoft AI Economy Institute, 2026). The breadth findings from Microsoft and the depth findings from Anthropic reinforce rather than replicate each other. The AI Matrix, published on SSRN in April 2025, predicted that lower-income countries would show characteristics consistent with Passive Dependency or Full Dependency, including directive interaction, coursework dominance, and concentrated use among technical early adopters, and that higher-income countries would show characteristics consistent with Elite Empowerment or Full Empowerment, including diversified use cases, iterative collaboration, and a broader distributional spread of users. This is precisely what

Anthropic, OpenAI, and Microsoft have independently found. The framework was not fitted to these findings. They arrived after it.

The paper proceeds as follows. Section 2 examines the capability overhang concept and the Anthropic adoption curve finding in more detail, identifying the specific analytical gaps that a two-dimensional framework is designed to fill. Section 3 presents the AI Matrix, grounds its two axes in established theory, and reinterprets capability overhang through its quadrant logic. Section 4 describes the AEI dataset, explains the operationalization choices, and addresses the limitations explicitly. Section 5 presents the three-stage typology and discusses the country-level evidence for each stage. Section 6 draws out the policy implications of a differentiated framework and addresses the path dependency problem that makes global convergence so much slower than within-country convergence. Section 7 concludes.

## **2. The Capability Overhang and Its Limits as a Policy Concept**

The notion that AI's most pressing problem is no longer capability but adoption marks a significant shift in how the technology's leading developers understand their own work. For most of the period since the release of GPT-3 in 2020, the dominant research question was whether AI systems were good enough, reliable enough, and safe enough to deploy at scale. By early 2026, OpenAI's public position was that this question had been substantially answered for a wide range of knowledge work tasks, and that the new frontier was behavioral and institutional rather than technical. Progress toward the full economic and social benefits of AI, the company argued, would depend as much on helping people use AI well as on advances in frontier models (OpenAI, 2026a). This reframing deserves to be taken seriously, not because OpenAI's commercial interests in expanding adoption are irrelevant, but because the empirical evidence from multiple independent sources supports the core observation. The gap between AI capability and AI use is real, it is large, and it is unevenly distributed in ways that have significant consequences for economic development and global inequality.

This section examines the capability overhang concept and the parallel findings from the Anthropic Economic Index in detail. It then situates both within the established literature on technology diffusion and the digital divide, where an analytically similar problem has an extensive history. It concludes by identifying three limitations of capability overhang as currently framed, limitations that the digital divide literature has already partially diagnosed, and that motivate the two-dimensional framework developed in Section 3.

## 2.1 The Capability Overhang: What the Evidence Shows

OpenAI's "Ending the Capability Overhang" report, released at the World Economic Forum in January 2026, presented evidence of a substantial gap between the most and least intensive users of AI tools. Power users, defined as those exploiting advanced reasoning, multi-step agentic workflows, and tool integration, relied on approximately seven times more advanced AI capabilities than typical users engaging in basic conversational interactions (OpenAI, 2026a). This gap was visible not only between individuals within the same organization or country, but between countries. Some nations showed substantially higher per-person use of advanced AI functions than their income levels would predict. Vietnam and Pakistan, for example, ranked among the highest users of agentic tools and advanced data analysis per person, ahead of many wealthier countries (OpenAI, 2026b). The implication was that capability overhang was not simply a function of access to AI infrastructure. Something else was shaping who used AI deeply and who used it superficially.

The Anthropic Economic Index (AEI) provides the most granular independent evidence for this conclusion. Across four reports published between January 2025 and January 2026, and based on privacy-preserving analysis of millions of Claude.ai conversations, Anthropic documented a consistent and striking geographic pattern. At the country level, a one percent increase in GDP per capita is associated with a 0.7 percent increase in Claude usage per capita (Appel et al., 2026a). But the relationship between income and AI adoption is not simply a matter of how much AI is used. It also shapes how AI is used. Higher-income countries show substantially higher rates of augmented interaction, in which users iterate collaboratively with Claude across multiple exchanges, learning, revising, and refining rather than issuing a single directive and accepting the result. Lower-income countries show substantially higher rates of directive interaction, in which users delegate a task and receive a response with minimal back-and-forth, and higher rates of coursework use relative to professional or personal applications (Appel et al., 2026a). The September 2025 report quantified this relationship precisely, finding that each one percent increase in population-adjusted Claude usage was associated with approximately a three percent reduction in automation rates after controlling for task mix (Appel et al., 2025). Higher adoption, in other words, is associated with more collaborative rather than more automated use, a finding that cuts against the intuition that AI diffusion naturally tends toward automation over time.

Data from Microsoft's AI Economy Institute confirms the same directional finding at the population level. In H2 2025, AI adoption reached 24.7 percent among the working-age population in the Global North against 14.1 percent in the Global South, with the gap widening rather than narrowing across the year (Microsoft AI Economy Institute, 2026). A use-case characterisation of Microsoft data, reported in Appel et al. (2026), further specifies that AI use for

educational purposes concentrates among lower per-capita income countries while AI use for leisure concentrates among wealthier populations. Across these programs, the depth and character of AI engagement varies systematically with income in ways that go well beyond simple differences in access to the technology.

The productivity implications of this variation are consequential. Anthropic's data shows that tasks requiring college-level engagement can accelerate task completion by a factor of 12 with AI assistance, compared to a nine fold acceleration for high-school-level tasks (Appel et al., 2026a). OpenAI's enterprise data shows that intensive AI users save substantially more time per day than casual users, with the most engaged workers reclaiming upwards of ten hours weekly (OpenAI, 2026a). If the depth of AI engagement is correlated with income at the country level, then the productivity gains from AI are not merely unevenly distributed in the present. They are compounding in ways that will widen existing gaps over time, as countries with more sophisticated AI use extract accelerating returns while those with shallower engagement capture only a fraction of the available benefit. This concern is consistent with projections from the International Monetary Fund, which has warned that high-income countries are positioned to capture the majority of AI-driven productivity gains, while lower-income countries risk deepening technological dependence (IMF, 2024).

## **2.2 An Established Precedent: The Digital Divide Literature**

The analytical problem that capability overhang describes is not new. A structurally similar challenge has been studied in the digital divide literature since the late 1990s, and the progression of that literature offers both conceptual resources and a cautionary tale for how the AI adoption gap is currently being framed.

Early digital divide research focused overwhelmingly on access, understood as physical connectivity to computers and the internet. The policy response was correspondingly focused on infrastructure provision: hardware distribution programs, connectivity subsidies, and the extension of broadband networks into underserved regions (Norris, 2001). This first-level divide framing assumed, implicitly, that the problem was one of reach, and that once populations could access digital technology, meaningful use would follow.

Research quickly complicated this picture. Hargittai (2002) demonstrated that substantial variation in the ability to use the internet effectively persisted even among populations with equivalent access, coining the concept of web-use skills as a distinct variable that the access-focused literature had obscured. Warschauer (2003) argued that the digital divide could not be understood

as a simple binary between the connected and the unconnected, and that technology access without the social, educational, and institutional supports for meaningful use produced little benefit. Van Dijk (2005) formalized this progression as a hierarchy of divides, distinguishing motivational access, physical access, skills access, and usage access as sequential barriers, each of which could be present or absent independently. Subsequent work developed this further, showing that differences in the quality and outcomes of internet use persisted well beyond the closing of simple connectivity gaps, and that these differences were systematically related to educational attainment, socioeconomic status, and institutional context (van Deursen & van Dijk, 2014; Selwyn, 2004).

The policy implication of this literature was clear, and largely ignored by the programs it was designed to inform. Infrastructure provision was necessary but not sufficient. Second-level barriers, which the digital divide literature variously called skills gaps, usage gaps, or outcome gaps, required different interventions from first-level connectivity barriers, and conflating the two produced programs that expanded access without improving outcomes for the populations most in need. Selwyn (2004, p. 347) made the point directly: the provision of access to technology is not the same as the provision of meaningful engagement with technology.

Thirty years later, the AI adoption literature is in danger of making the same analytical error. The capability overhang concept, as currently deployed, is a first-level framing of what is demonstrably a multi-level problem. The evidence from Anthropic, OpenAI, and Microsoft consistently shows that the most important variation in AI use across countries is not whether populations have access to AI tools, but how they engage with those tools when they do have access. This is precisely the second-level divide that Hargittai, Warschauer, and van Dijk documented for the internet. The difference is that the productivity stakes of second-level AI divides are substantially higher than those of second-level internet divides, because the potential productivity gains from sophisticated AI use, measured in multiples of task completion time, are larger than anything the internet access literature was working with.

The AI readiness indices that currently dominate policy assessment of AI adoption have replicated the first-level bias of early digital divide research. The Stanford AI Index (Stanford HAI, 2024), the Oxford Insights Government AI Readiness Index (Oxford Insights, 2024), and the Tortoise Global AI Index (Tortoise Media, 2023) all evaluate countries primarily on the basis of infrastructure, investment, research output, and regulatory frameworks. They measure the conditions under which AI can be accessed, not the conditions under which it can be used effectively. As a result, they provide limited guidance for the question that the capability overhang literature has made

central: why do countries at similar income levels show dramatically different profiles of AI engagement, and what would actually change that?

### **2.3 Three Analytical Limitations of the Capability Overhang as Currently Framed**

Against this background, three specific limitations of the capability overhang concept constrain its usefulness as a guide to policy, each of which mirrors a limitation that the digital divide literature has already identified and partially addressed.

The first limitation is that the concept is undimensional. Capability overhang describes a gap between AI's potential and its realization without specifying what kind of gap it is. When OpenAI observes that some countries use AI seven times more intensively than others, and that this difference is not fully explained by income, the obvious next question is what does explain it. The capability overhang framework does not answer this question because it treats the gap as a single phenomenon. The digital divide literature resolved an equivalent ambiguity by distinguishing access from skills from outcomes as separate analytical dimensions (van Dijk, 2005). A corresponding distinction for AI, between the structural availability of tools on the one hand and the capacity to engage with them productively on the other, is necessary for the same reason: without it, the gap is identified but not diagnosed, and prescription without diagnosis produces misdirected policy.

The second limitation is that the framework is largely atheoretical. The capability overhang literature documents patterns without providing a mechanism. Why does the gap persist even as AI tools become cheaper and more accessible? Why do some countries at similar income levels show dramatically different engagement profiles? Why have global gaps between countries remained stable even as within-country gaps in the United States are narrowing rapidly, at a pace Anthropoc estimates is roughly ten times faster than the diffusion of comparable technologies in the twentieth century (Appel et al., 2026a)? Technology diffusion research has well-developed theoretical resources for addressing these questions, including Rogers's (2003) diffusion of innovations framework, which identifies social system characteristics, communication channels, and adopter heterogeneity as determinants of diffusion speed and pattern. Path dependency theory, developed by David (1985) and Arthur (1989), provides a further mechanism for understanding why gaps persist. Early institutional choices and resource allocations create self-reinforcing trajectories that make alternative paths progressively harder to access. Neither framework has been applied systematically to the AI adoption gap, and the capability overhang literature does not fill this theoretical vacancy.

The third limitation follows from the second. Because the capability overhang concept lacks a theoretical account of why gaps persist, it also lacks a framework for understanding what would change them. The policy recommendations flowing from the Davos discussion, and from OpenAI's for Countries initiative, center on expanding access through infrastructure investment, building AI skills through education programs, and creating institutional readiness through governance reform. These are reasonable starting points, and they parallel the infrastructure-plus-skills approach that the digital divide literature eventually converged on after two decades of access-only policy. But they are not differentiated to the specific situation of countries at different positions on the adoption curve. A country whose primary constraint is infrastructure will not be helped much by AI literacy programs targeted at populations who cannot yet reach the tools. A country whose professional class already uses AI intensively, but where that use has not diffused beyond a narrow technical elite, will not be helped by additional infrastructure investment reaching populations who lack the capacity to exploit it productively. The capability overhang concept, as currently formulated, does not provide the analytical tools to distinguish these cases from each other, or to prescribe accordingly.

## **2.4 Access and Agency as Distinct Dimensions**

Addressing these three limitations requires a framework that treats the gap between AI capability and AI use as multidimensional, specifies the mechanisms that maintain it, and generates differentiated predictions about what kinds of intervention would be effective in different contexts. Building on the conceptual architecture developed in the digital divide literature, we propose that the critical distinction is between access and agency, understood as genuinely separate dimensions rather than points on a single spectrum.

Access, as we use the term, refers to the structural conditions that determine whether a given population can reach AI tools at all. It encompasses physical and digital infrastructure, the affordability of devices and connectivity, the availability of AI services in relevant languages, and the regulatory and political conditions that either enable or restrict access. Access is primarily a function of material and institutional resources, and it responds in principle to investment in those resources. The relationship between income and access is strong and well-documented across digital technologies generally (Norris, 2001; van Dijk, 2005), AI included, which is why GDP per capita explains so much of the variation in AI adoption rates at the country level.

Agency, as we use the term, refers to the capacity to engage with AI tools in ways that generate productive outcomes. It encompasses the sophistication of prompting and interaction, the willingness to iterate across multiple exchanges rather than accepting a single output, the ability to evaluate AI-generated content critically rather than passively, and the educational and

professional dispositions that support these behaviors. Agency is a function of human capital, educational systems, and the institutional contexts in which AI is used. It is not simply a matter of knowing that AI tools are available, or having been given access to them. This usage-level dimension is precisely what Hargittai (2002) identified as the second-level digital divide, what van Deursen and van Dijk (2014) later called the usage gap, and what Warschauer (2003) described as the social embeddedness of meaningful technology engagement. We apply the same analytical logic to AI, with the additional observation that the productivity consequences of the agency gap are considerably larger than those the internet-access literature was working with.

The distinction matters empirically because the two dimensions are not perfectly correlated. The OpenAI finding that Vietnam and Pakistan rank highly on agentic AI use despite moderate income levels illustrates that agency can develop without broad access having been achieved first, typically among technically skilled professional populations. The Anthropic finding that India ranks second globally in absolute Claude usage but 101st out of 116 countries in per-capita usage, with more than half of all use concentrated in four IT-sector states (Appel et al., 2026b), illustrates the complementary case: a country where access and high agency coexist for a professional minority while the broader population remains largely unreached. In the Indian case, the binding constraint for most of the population is not access in the infrastructure sense. It is the distributional failure through which the benefits of existing high-agency use are not being spread. Critically, neither the capability overhang concept nor the AI readiness indices currently in use provide the tools to characterize this situation or to prescribe for it. A note on operationalization follows from this distinction: the empirical proxy for access used in Section 4 captures effective reach or realized access in practice rather than structural access in the full theoretical sense defined here. That gap between construct and proxy is acknowledged explicitly in Section 4.3.

Section 3 introduces the AI Matrix as a framework that does. By organizing countries along the access and agency dimensions simultaneously, and specifying the transition mechanisms through which countries move between quadrants, the AI Matrix provides the analytical structure that the capability overhang literature currently lacks and that the digital divide literature arrived at only after decades of first-level policy failure.

### **3. The AI Matrix: A Framework for Two-Dimensional Decomposition**

The argument in Section 2 establishes that capability overhang is undimensional as currently framed, atheoretical in its account of why gaps persist, and transitionally agnostic about what would actually close them. Filling these three gaps requires a framework that organizes countries along both dimensions simultaneously, grounds their positions in an account of mechanism rather than correlation, and specifies the conditions under which movement between positions becomes

possible. The AI Matrix (Simpson, 2025a) provides this structure. This section grounds its two axes in the theoretical traditions on which it draws, and reinterprets capability overhang through its quadrant logic.

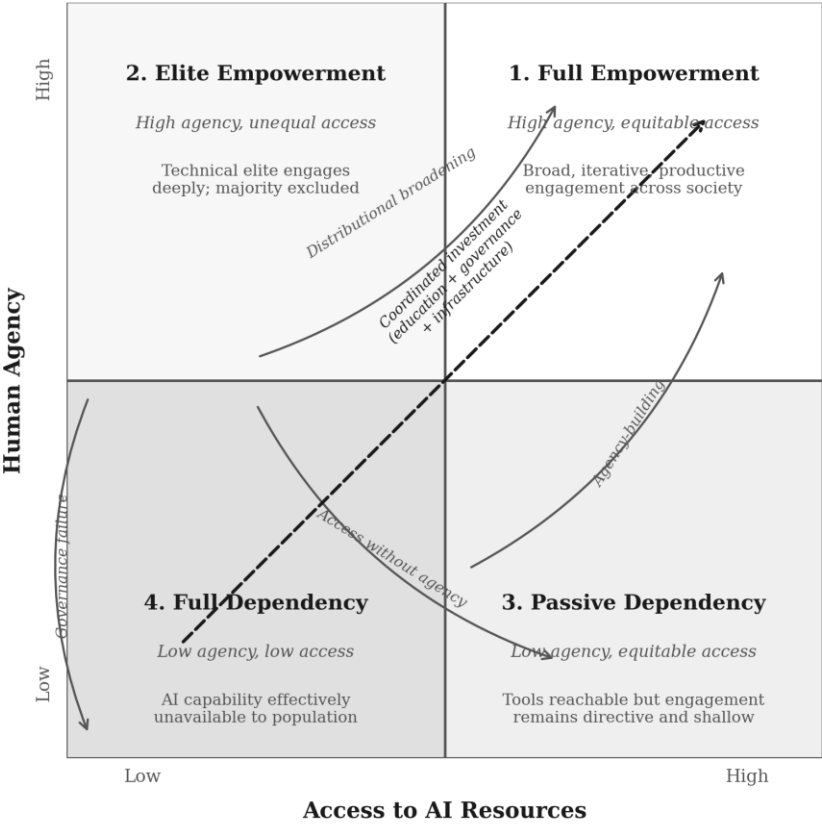
### **3.1 The Framework**

The AI Matrix organizes actors along two independent axes. The horizontal axis measures access to AI resources, understood as the structural conditions that determine whether a given population can reach AI tools at all. This encompasses physical and digital infrastructure, the affordability of devices and data, the availability of AI services in relevant languages, and the regulatory and political environment that either enables or restricts use. The vertical axis measures human agency, understood as the capacity to engage with those tools in ways that generate productive outcomes. This encompasses the sophistication of interaction, the willingness to iterate and refine rather than simply delegate, and the educational and institutional conditions that support critical rather than passive engagement.

The intersection of these two axes produces four quadrants, each describing a qualitatively distinct relationship with AI technology. Full Empowerment, in the upper right, characterizes actors with both broad access and high agency, populations that can reach AI tools and use them in ways that generate genuine augmentation of human capability, fostering innovation, inclusion, and sustainable growth. Elite Empowerment, in the upper left, characterizes actors with high agency but unequal or concentrated access: a technically skilled professional minority that engages deeply with AI while the broader population remains largely unreached, deepening inequality even as it produces innovation. Passive Dependency, in the lower right, characterizes actors with broad access but low agency: populations that can reach AI tools but engage with them directly and shallowly, producing reliance on automated systems while reducing the scope for critical judgment and creative engagement. Full Dependency, in the lower left, characterizes actors with both low access and low agency: populations for whom AI capability is, in practical terms, unavailable, and who face exclusion from AI-driven opportunities entirely.

The framework's most important design feature is that it is explicitly dynamic. Quadrant positions are not fixed assessments of national character or economic destiny. They are snapshots of a trajectory, and the framework specifies the conditions and interventions through which actors move between quadrants. These transitions are neither automatic nor inevitable. They require intentional, sustained effort across multiple dimensions, and the direction of movement is not predetermined. A country in Full Dependency can move toward Elite Empowerment if access expands among technically skilled populations before broader diffusion occurs, or toward Passive Dependency if access expands without concurrent agency-building. A country in Elite

Empowerment can move toward Full Empowerment through deliberate distributional broadening, or slide toward greater concentration if the professional minority retains its advantages while the majority falls further behind. This transition logic is mapped explicitly in Figure 1 below.



**Figure 1. The AI Matrix: Empowerment or Dependency?**

Note: Arrows indicate transition pathways between quadrant positions. Movement is neither automatic nor predetermined; each pathway requires specific institutional conditions and interventions described in Section 3.2. The dashed diagonal represents the coordinated investment trajectory (education, governance, and infrastructure) required to move directly from Full Dependency to Full Empowerment.

Figure 1 presents the full conceptual architecture of the AI Matrix as used in this paper. It should be read alongside the quadrant descriptions above and the theoretical grounding in Section 3.2. The transition arrows are the framework's primary contribution beyond existing readiness indices: they specify not only where countries sit but the mechanisms and conditions under which they move. The empirical counterpart of this conceptual diagram, showing where the 172 countries in the AEI dataset actually fall, appears as Figure 2a in Section 4.3.

### **3.2 Theoretical Grounding**

The AI Matrix draws on four established theoretical traditions. They do not all bear equal argumentative weight. Path dependency and governance theory are the load-bearing mechanisms, explaining why gaps persist and what institutional conditions would close them. Socio-economic resilience and critical theory function as interpretive and normative frames, clarifying what is at stake in the distributional outcomes the framework describes without generating independent empirical predictions. Together, the four traditions provide the mechanistic account that the capability overhang literature currently lacks and that the digital divide literature only partially supplied.

The first tradition is path dependency theory, developed by David (1985) and Arthur (1989). Path dependency holds that early choices among competing technological or institutional options, once made, become progressively self-reinforcing through increasing returns, sunk costs, and network effects, making alternative trajectories harder to access over time. Applied to AI adoption, path dependency explains why Full Dependency is the most stable of the four quadrant positions. The institutional conditions that support agency-building, capable educational systems, governance frameworks that reward critical knowledge work, and economic structures that create demand for higher-order skills, are themselves products of prior developmental investment. Countries that lack these conditions cannot build agency simply by gaining access to AI tools, because the foundations on which agency development depends are not yet in place. The dominance of large corporations and high-income countries in AI research and development compounds this dynamic, creating barriers that consolidate technological dependency in underserved regions and making the concentration of capability self-reinforcing over time (Arthur, 1989; Acemoglu & Restrepo, 2019). Path dependency does not imply permanence: the developmental trajectories of Singapore and South Korea demonstrate that self-reinforcing cycles can be interrupted through sustained, coordinated investment. But interrupting them requires far more than access provision alone.

The second tradition is socio-economic resilience theory, which examines how systems adapt to shocks and disruptions and maintain the capacity to recover, reorganize, and innovate (Holling, 1973; Folke et al., 2010). Originally developed in ecology, the resilience framework has been extended to economic and social systems, where it identifies the conditions under which communities can withstand technological disruptions while preserving human agency and equitable opportunity. In the AI Matrix, resilience theory provides the mechanism for understanding why Elite Empowerment is an inherently unstable quadrant position that can move in either direction. Nations with diversified economies and adaptive education systems, such as Singapore, demonstrate greater capacity to move toward Full Empowerment because their

institutional resilience allows agency-building to diffuse across the population rather than remaining concentrated in a technical elite (Folke et al., 2010). By contrast, economies where AI adoption is not accompanied by investments in social safety nets, workforce development, and sectors less susceptible to automation face systemic vulnerabilities as displacement accumulates and the distributional gap between AI-engaged and AI-excluded populations widens. The Indian case illustrates this instability precisely: a country where the IT sector's sophisticated engagement generates enormous aggregate productivity while most of the population remains below the agency threshold, and where the gap is widening rather than closing (Appel et al., 2026a).

The third tradition is governance theory, particularly the multi-level governance framework developed by Stoker (1998) and the polycentric governance approach articulated by Ostrom (2009). These frameworks emphasize that institutional arrangements, rather than market forces alone, are the primary determinants of how technology is deployed and who benefits from it. In the AI Matrix, governance theory provides the mechanism for the most important empirical asymmetry in the AEI data: the sharp divergence between within-country convergence and between-country divergence in AI adoption. Within the United States, Anthropic estimates that per-capita AI usage would equalize across states within two to five years, roughly ten times faster than comparable technology diffusions in the twentieth century (Appel et al., 2026a). Globally, gaps between countries show no evidence of convergence. The governance explanation is that within a common regulatory and institutional framework, above the access threshold, diffusion is driven primarily by workforce composition and institutional learning, both of which operate relatively rapidly when the foundational conditions are present. Across countries, the access threshold itself is a governance product, shaped by political choices about infrastructure investment, trade policy, and regulatory openness. Closing it requires international coordination that currently does not exist at the scale the problem demands. Where governance frameworks prioritize market efficiency over equity, as Jessop (2002) documents in the context of neoliberal governance more broadly, the concentration of AI capability among a small number of corporations and nations tends to be self-reinforcing rather than self-correcting. Supplementing Stoker and Ostrom with the institutional economics literature, Acemoglu, Johnson, and Robinson (2005) provide the most technically rigorous account of why governance frameworks are not merely enabling conditions for technology diffusion but fundamental determinants of long-run distributional outcomes.

The fourth tradition is critical theory, which provides a lens for understanding the socio-political implications of AI adoption and the need to challenge the assumption that technology diffusion is inherently progressive (Feenberg, 1991). From this perspective, movement toward Full Empowerment requires deliberate efforts to democratize AI development, ensuring that

marginalized voices shape the trajectory of adoption rather than simply receiving its outputs. The dependency scenarios in the AI Matrix are not merely technical failures: they reflect choices about whose interests AI systems serve and whose they do not (Eubanks, 2018). For the empirical analysis in this paper, critical theory operates primarily as a background orientation, informing why the agency dimension of the framework is treated as a political and institutional phenomenon rather than a purely technical one, and why the distributional question of who benefits from high-agency AI use is treated as analytically central rather than secondary to the question of aggregate adoption rates.

### **3.3 Capability Overhang Through the AI Matrix Lens**

The framework's most important theoretical contribution is its reinterpretation of capability overhang as three structurally distinct problems rather than one. The policy implications of this reinterpretation are developed fully in Section 6. The analytical point is established here.

In Full Dependency, capability overhang is fundamentally an access problem. The gap between what AI can do and what a given population does with it exists primarily because that population cannot reach the tools. The framework implies that the appropriate intervention is access expansion: infrastructure, affordability, language support, and regulatory openness. This is a differentiated policy hypothesis grounded in the theoretical logic of the quadrant positions rather than a prescription demonstrated by this evidence. The framework also implies that access expansion alone is not sufficient. Without concurrent investment in the conditions that support agency development, access expansion moves countries from Full Dependency to Passive Dependency rather than toward Full Empowerment. The tools become reachable without becoming productive.

In Passive Dependency, capability overhang is fundamentally an agency problem. Access is present but effective engagement is not. Users reach the tools but interact with them directly and shallowly, capturing a fraction of the available benefit. The education-first adoption pattern that the AEI data identifies in Indonesia, Peru, Ecuador, and several sub-Saharan African countries is consistent with this position: high coursework use, directive interaction patterns, and task success rates that suggest users are receiving outputs without building the iterative engagement that generates compounding returns. The framework implies that the appropriate intervention is not more access provision but agency-building: curriculum redesign, AI literacy programs focused on iterative rather than passive use, and professional development frameworks that treat prompting as a design skill rather than a search behavior. This inference applies to the population already above the access threshold visible in the data; it does not speak to populations still absent from it.

In Elite Empowerment, capability overhang is a distributional problem. The technology is well exploited by an identifiable professional minority and chronically underused by the majority. The aggregate figures for a country in Elite Empowerment may look impressive, and they often are: India's second-place ranking in absolute Claude usage reflects genuine, sophisticated engagement by a large and highly skilled IT sector. But per-capita usage that ranks 101st out of 116 countries (Appel et al., 2026b), with more than half of all interactions concentrated in four states, tells a different story. The capability overhang in this case is not a failure of access or a failure of agency among those who have access. It is a failure of distribution: the conditions that produced high-agency use in one segment of the population have not been extended to the rest. The framework implies that the appropriate intervention is structural broadening: occupational diversification beyond the technical sector, geographic extension of digital infrastructure and AI literacy, and economic policies that create demand for higher-order AI engagement across a wider range of industries.

This three-way decomposition is the theoretical foundation for the empirical analysis in Sections 4 and 5. It offers a theoretically grounded account consistent with why the adoption curve documented independently by Anthropic and OpenAI takes the shape that it does, why global gaps between countries have remained stable while within-country convergence accelerates in wealthier nations, and why undifferentiated access-expansion policies have so far failed to close the engagement gap. These interpretations remain to be tested with causal rigour; they are offered here as theoretically motivated expectations consistent with the observed patterns. The following section describes the data and operationalization choices through which we test whether the distribution of countries in the AEI dataset is consistent with these predictions.

#### **4. Data, Operationalization, and Limitations**

The empirical analysis in this paper rests on a single primary dataset supplemented by external access indicators. This section describes both in detail, explains the operationalization choices through which we translate raw variables into proxies for the AI Matrix's two dimensions, and sets out the limitations of this approach with precision. The goal is not to minimize those limitations but to establish exactly what the data can and cannot support, so that the interpretive claims in Section 5 are appropriately bounded.

## 4.1 The Anthropic Economic Index Dataset

The Anthropic Economic Index is a program of ongoing empirical research into AI's effects on labor markets and the economy, conducted through a privacy-preserving analysis of Claude.ai conversations. The dataset used in this paper, released publicly by Anthropic in January 2026 with an explicit invitation to external researchers, covers 818,492 conversations from Claude.ai Free and Pro users across 173 countries during November 13-20, 2025. The data reflects primarily Claude Sonnet 4.5 usage. It is the most geographically comprehensive observed dataset of AI interaction behavior currently publicly available to researchers.

The dataset is structured at the conversation level. For each conversation, Anthropic's Clio system, a privacy-preserving analysis tool that processes anonymized transcripts without retaining identifiable content, estimated five categories of variable. Use case classifies each conversation as primarily work-related, academic coursework, or personal. Collaboration pattern assigns each conversation to one of five categories: directive, in which the user delegates a task to Claude and accepts the output; task iteration, in which the user refines and develops a task across multiple exchanges; learning, in which the user asks Claude to explain or teach; feedback loop, in which the user reviews Claude's output and requests revisions; and validation, in which the user asks Claude to verify their own work. Task success records Claude's self-assessment of whether it completed the task as requested. Education level provides an estimate, in years of formal schooling on the standard US scale, of the complexity of both the user's prompt and Claude's response, where 12 years corresponds to a high school diploma and 16 years to a bachelor's degree. AI autonomy scores each conversation from 1 to 5 based on the degree to which the user delegated decision-making and execution to Claude, where 1 indicates a human-led interaction with minimal AI contribution and 5 indicates fully autonomous AI execution.

At the global level, the dataset shows that 46.0% of conversations were work-related, 34.7% personal, and 19.3% academic coursework. The most common collaboration pattern was directive at 31.7%, followed by task iteration at 27.2%, learning at 19.7%, feedback loop at 13.6%, and validation at 4.8%. The global task success rate was 66.9%. The average education level of prompts was 12.2 years, closely matched by the average education level of responses, reflecting Claude's tendency to calibrate response complexity to the sophistication of the input. The mean autonomy score was 3.4 out of 5, with score 4 the most common at 48.6% of conversations. On time savings, the mean estimated task time without Claude was 3.1 hours against a mean conversation length of 15 minutes, implying a roughly 12-times speedup on a mean-to-mean basis and approximately 8 times on a median-to-median basis. After adjustment for the 66.9% task success rate, Anthropic's own productivity estimate falls to an annual US labor productivity uplift of 1.0 to 1.2 percentage points.

## 4.2 Operationalizing Agency

The agency dimension of the AI Matrix refers to the capacity to engage with AI tools in ways that generate productive outcomes. The AEI variables do not measure this directly, but several of them function as observable proxies for the behavioral patterns that high agency, as the framework defines it, would be expected to produce.

We construct a country-level agency composite from four components. The first is the proportion of conversations classified as active co-creation, defined as task iteration, learning, and validation combined. These patterns require the user to engage iteratively with Claude, either refining an output, building understanding, or exercising critical judgment over AI-generated content. Higher rates indicate more agentic engagement. The second component is the proportion classified as directive or no identifiable pattern, which captures passive delegation and shallow engagement. Higher rates on this component indicate lower agency. The third is the task success rate, which serves as an outcome marker: higher success rates, particularly on work-related tasks, are consistent with users who can specify tasks clearly, evaluate outputs critically, and iterate when the result falls short. The fourth is the education level of prompts, which serves as a sophistication proxy: prompts requiring more years of formal education to write tend to reflect more precise task specification and more complex underlying needs.

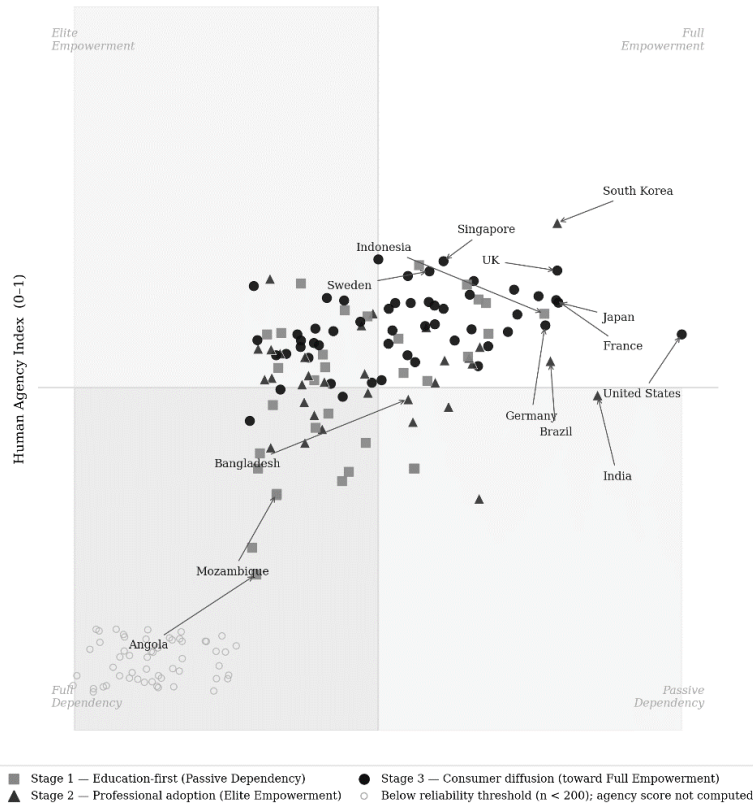
We treat the agency composite as an ordinal indicator rather than a cardinal measure. Countries with higher active co-creation rates, lower directive rates, higher task success rates, and higher prompt education levels are coded as exhibiting interaction patterns consistent with higher agency in the AI Matrix sense. We are explicit that this measures observed behavior, not agency as a psychological or institutional property. It is a proxy for what the framework conceptualizes, constructed from the variables available in the data, and should be interpreted accordingly. The stage assignment rules, which translate these proxies and raw variables into the three-stage typology, are stated here for transparency. Countries are assigned using a sequential rule applied to the raw AEI variables: Stage 1 (Education-first) if coursework share exceeds 30%; Stage 2 (Professional adoption) if work share exceeds 48% and coursework share is below 25%; Stage 3 (Consumer diffusion) if personal use share exceeds 38%. Countries not assigned by these primary rules are classified by agency composite score: 0.55 or above = Stage 3; 0.42 to 0.55 = Stage 2; below 0.42 = Stage 1. These thresholds operationalize the qualitative distinctions described in Section 5; they are not optimised statistical cutpoints. Modest variation in the thresholds does not materially alter the core country groupings, as reported in Section 6.3.

### 4.3 Operationalizing Access: Raw Volume

The AEI dataset cannot provide a direct measure of the access dimension. It covers only countries where Claude is already used in sufficient volume to appear in the data, which means it is, by construction, a sample of countries above some minimum access threshold. For the horizontal axis of Figure 2a, we use the log base-10 of each country's raw conversation count as an access proxy, min-max normalised across all 172 countries in the dataset. Log transformation is applied to compress the extreme right-skew of the raw distribution, in which the United States accounts for approximately 22 percent of all conversations. This measure captures the actual scale of AI engagement as a global economic fact. The United States generating 219,700 conversations in one week, India generating 58,100, and Tanzania generating 292 are not methodological artefacts: they reflect the real distribution of AI-generated productive activity across the world. This concentration is a central finding of the paper, not a feature to be normalised away. A note on terminology is warranted here. Throughout the paper we use "access" as shorthand for what is more precisely described as effective reach or realized access in practice: the observed extent to which a population is engaging with AI tools in this dataset. This measure conflates structural access (the availability of infrastructure, affordability, and regulatory conditions) with adoption behavior (the decision to use tools that are available). The two cannot be cleanly separated within the AEI data alone, and readers should interpret all access-axis positions as measures of effective reach rather than structural access in the AI Matrix's theoretical sense.

We note that the Oxford Insights Government AI Readiness Index (Oxford Insights, 2024) is precisely the kind of static, snapshot-oriented readiness assessment that the AI Matrix is designed to supplement and critique. We use the AEI volume measure rather than readiness indices because it reflects observed behavior rather than assessed potential, and because its limitations are transparent and consistent across countries. It does not imply endorsement of raw volume as the theoretically ideal access measure. Section 4.3a below develops a per-capita complement.

**Figure 2a. Country Distribution in the AI Matrix: Raw Conversation Volume**  
**AEI November 2025 data; all 172 countries with any conversation volume**



Access to AI Resources (0-1)

Derived from  $\log_{10}$  conversation volume; min-max normalised across all 172 countries. Full derivation: Appendix Table A1.

Note: Countries absent from the dataset entirely (~23) are plotted in the Full Dependency quadrant. Their absence reflects access failure rather than non-adoption.

Low-confidence countries ( $n < 200$ , hollow circles) are plotted at their access position with agency treated as indeterminate. Selected countries labelled for orientation.

Approximately 23 further countries are absent from the dataset entirely — predominantly low-income states in sub-Saharan Africa, Central Asia, and the Pacific. Full Dependency is the modal condition for most of the global south.

Figure 2a shows where AI engagement is concentrated in absolute terms across the 172 countries in the dataset. The rightward clustering of the US, India, and several other large economies reflects the scale of their contribution to global AI activity in November 2025. This concentration is one of the paper's central empirical observations. Figure 2b immediately below reframes the same data on a per-capita basis, revealing a different distributional picture in which population size is held constant. The two figures should be read together. Figure 2a answers where AI economic activity is happening; Figure 2b answers how deeply it has penetrated each country's population.

### 4.3a Raw Volume and Per-Capita Access: Two Distinct Findings

The raw conversation volume measure used for the access axis in Figure 2a and the per-capita measure introduced in Figure 2b are not simply two ways of calculating the same thing. They measure different phenomena, both of which are analytically important, and the differences between them are themselves findings rather than methodological noise.

Raw volume captures the scale of AI engagement as a global economic and geopolitical fact. The United States generated approximately 219,700 conversations in the November 2025 week, representing 22 percent of all observations in the dataset. India generated 58,100, approximately 5.8 percent. Together these two countries account for more than a quarter of all recorded AI interactions in what is the most geographically comprehensive observed dataset of AI use currently available to researchers. This concentration is not an artefact to be divided away. It reflects the actual distribution of AI-generated economic activity across the world in 2025, and it has direct implications for questions of productivity, labor market effects, and the global distribution of AI-driven value creation. When we say the capability overhang matters, the volume figures are part of what we mean: the absolute scale of productive AI engagement is heavily concentrated in a small number of high-income countries and one large middle-income one.

Per-capita volume reframes the same data as a population-level reach question. On this measure, Singapore leads all countries in the dataset at 863 conversations per million population, followed by Israel at 817, the United States at 656, Canada at 597, South Korea at 592, and Australia and New Zealand both above 550. The Nordic countries, which sit prominently in the upper right of Figure 2a, remain in the upper half of the per-capita distribution at between 380 and 480 per million. What changes most dramatically is India. Second in raw volume, India falls to approximately 41 conversations per million population, below Indonesia (91), Brazil (128), and Mexico (177). Pakistan falls to 39 per million. These figures make the Elite Empowerment characterisation of India and Pakistan more visually precise than the raw volume chart can achieve: the high aggregate counts reflect population scale and IT sector concentration, not population-level reach. The per-capita chart shows what the raw volume chart implies but does not directly display.

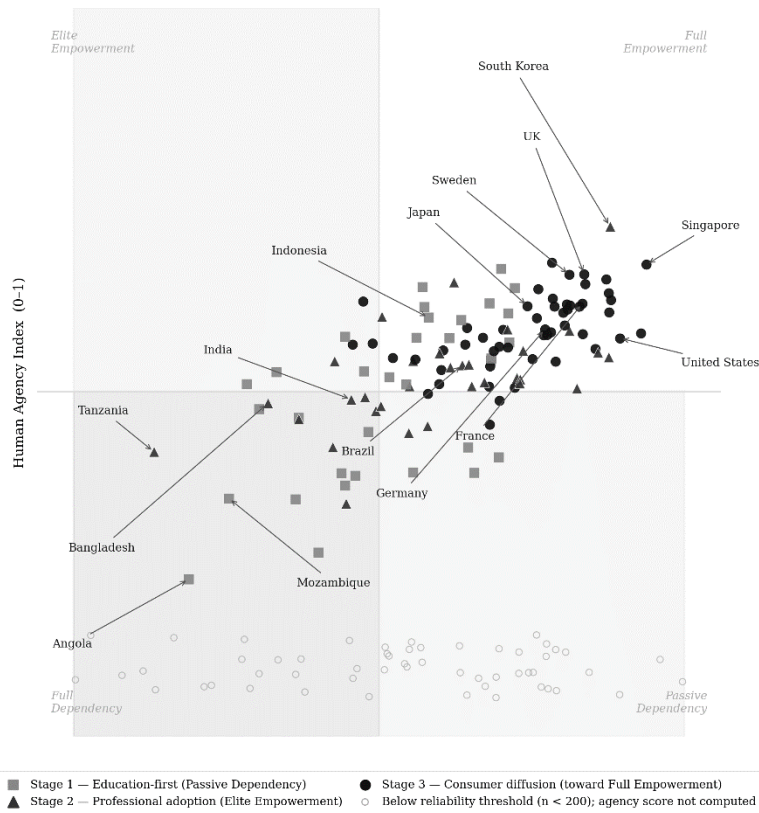
The bottom of the per-capita distribution is as instructive as the top. Tanzania records 5 conversations per million, Angola 8, Mozambique 11, Uganda 14. These are countries present in the dataset, meaning they have crossed the minimum access threshold, but whose populations are reached at a rate that is more than 100 times lower than Singapore's and roughly 130 times lower than the United States. The absolute gap that the raw volume figure conveys as a single

dominant mass in the upper right of Figure 2a is rendered in per-capita terms as a gradient of extraordinary range: more than two orders of magnitude between the most and least reached populations among countries that appear in the data at all. Countries absent from the dataset entirely, as noted in both figures, represent a further layer of exclusion below even these low per-capita rates.

One finding is specific to the per-capita analysis and not visible in the raw volume measure. South Korea's per-capita score of 592 per million places it among the highest-reach countries in the dataset, comparable to Australia and Canada, despite being classified as Stage 2 on interaction patterns. This pairing of high population-level reach with professional rather than consumer diffusion patterns is exactly what the transition logic of the AI Matrix predicts for a country approaching the Stage 2 to Stage 3 boundary: access has become broadly distributed before the interaction pattern has shifted. The per-capita measure makes this dynamic visible in a way the raw volume figure cannot.

The two figures should therefore be read together rather than as alternatives. Figure 2a asks where AI engagement is concentrated in absolute terms, which matters for understanding the present distribution of AI-generated productivity and the scale of the global capability overhang. Figure 2b asks how deeply AI has penetrated each country's population relative to its size, which matters for understanding the structural reach of AI adoption and the realistic prospects for diffusion. Both questions are analytically necessary. Neither answer displaces the other.

**Figure 2b. Country Distribution in the AI Matrix: Per-Capita Access**  
 AEI November 2025 data; conversations per million population, all 171 matched countries



Access to AI Resources (0-1)  
 Derived from  $\log_{10}$  conversations per million population; min-max normalised across 171 countries. Population data: World Bank/UN 2024 estimates.

Note: The agency axis and stage assignments are identical to Figure 2a. Only the horizontal (access) axis changes. Singapore leads at 863 conversations per million; India falls to 41 per million despite ranking 2nd in raw volume.

~12 countries present in Figure 2a are absent here (no reliable population denominator). Full Dependency remains the modal condition for most of the global south. See Section 4.3a.

Figure 2b makes visible the population-level reach dimension that raw volume cannot show. India's dramatic leftward shift relative to Figure 2a is the most analytically important feature: it converts the Elite Empowerment characterisation from a textual argument to a directly visible positioning. A country with 58,000 conversations in one week sounds like a major AI adopter; 41 conversations per million population against Singapore's 863 makes the distributional failure precise. South Korea's high per-capita position combined with its Stage 2 interaction profile is also newly visible here, reinforcing the transition argument made in Section 5.3.

#### 4.4 Data Utility Review

The combination of the AEI interaction data and external access indicators allows us to place countries in approximate positions within the AI Matrix's two-dimensional space and to identify the clustering patterns described in Section 5. It does not allow us to do several things that a stronger research design would permit, and these constraints require explicit acknowledgment.

The dataset is cross-sectional. It covers one week of conversations on one platform, which means it cannot establish whether the patterns observed are stable over time, whether they reflect long-run adoption trajectories or short-run fluctuations, or how they would appear on a different platform with a different user base. The present analysis uses the November 2025 wave, which is the release publicly available in sufficient detail for country-level analysis.

The dataset cannot establish causal direction. The correlation between income and interaction sophistication, between GDP per capita and active co-creation rates, is consistent with the AI Matrix's prediction that higher-income countries tend toward Full Empowerment or Elite Empowerment. It is also consistent with alternative explanations, including selection effects in which early Claude adopters in lower-income countries are disproportionately students, or platform effects in which Claude's design elicits more directive interaction from users with less prior experience of AI tools. We cannot distinguish between these explanations in a cross-sectional, single-provider dataset.

The invisible countries present the most significant analytical constraint. Countries absent from the data, or present in very small numbers below the reliability threshold of approximately 200 conversations, are almost certainly those closest to Full Dependency in the AI Matrix sense. The framework's predictions about Full Dependency cannot be tested against this data because the populations most affected by access failure are precisely those not represented. The empirical analysis in Section 5 should be understood as within-frontier variation, describing the distribution of countries that are already above the access threshold, rather than a complete global picture.

Finally, the task success variable is a self-assessment by Claude rather than an independent evaluation, and the collaboration pattern classification is a product of Clio's automated analysis rather than human coding. Both are subject to misclassification, and the rates of misclassification are not known. The education level estimates are likewise Claude's own inference from conversation content rather than verified demographic data about users. These measurement uncertainties do not undermine the broad patterns the data reveals, but they set a ceiling on the precision with which individual country positions can be characterized.

Within these constraints, the patterns in the data are consistent with the AI Matrix's predictions and sufficient to support the three-stage typology developed in Section 5. The consistency between the AEI findings and OpenAI's independently collected evidence provides partial external validation that the patterns are not artifacts of a single platform. Microsoft's AI Economy Institute data provides independent corroboration from a different vantage point: where the AEI captures depth of engagement among existing users, the Microsoft data captures breadth of adoption across populations (Microsoft AI Economy Institute, 2026). The two measures are methodologically distinct and their directional consistency strengthens confidence in the overall picture. A systematic cross-provider comparison using matched country samples would be a considerably stronger test, and we identify it as the most important direction for future research.

## **5. Empirical Patterns: A Three-Stage Adoption Typology**

### **5.1 Overview**

A preliminary note on the data is necessary before describing the stages. The AEI dataset covers Claude.ai Free and Pro consumer product conversations exclusively. It does not include API usage or enterprise deployment. What the stages therefore describe is the trajectory of AI adoption as experienced by individual users engaging with a general-purpose consumer AI tool. The data capture the diffusion of AI into personal and professional life at the individual level, not the integration of AI into organisational workflows at the institutional level. This distinction matters for interpretation: a country with concentrated enterprise API deployment but limited consumer use could appear at a low stage in this typology while being substantively advanced in institutional adoption. The framework as operationalised here cannot speak to that dimension.

The distribution of countries in the AEI dataset, when organized along the agency composite constructed in Section 4 and situated against external access indicators, does not produce a uniform continuum. It produces three regions of greater density that correspond to qualitatively distinct patterns of AI use. We characterize these as the education-first stage, the professional adoption stage, and the consumer diffusion stage.


These three stages represent a progression in the structural complexity of adoption, not a ranking of individual interaction quality. The complexity gradient operates across three observable dimensions. First, use case breadth: Stage 1 use is concentrated in a single domain (academic coursework), Stage 2 use is concentrated in a different single domain (professional work), and Stage 3 use is distributed across work, personal, and creative purposes simultaneously. Breadth of use is a reasonable proxy for the degree to which AI has become a general-purpose tool rather

than a specialist one. Second, interaction passivity: directive rates, defined as the proportion of conversations in which a user delegates a task and accepts the output without further engagement fall monotonically across stages. The stage-level means are 37.1% in Stage 1, 32.6% in Stage 2, and 28.6% in Stage 3. Lower directive rates indicate that users are more likely to engage iteratively, critically, or evaluatively rather than treating the tool as a one-shot output machine. It is worth noting that raw task iteration rates are approximately flat across stages at around 27%, which means the complexity gradient is not captured by iteration frequency alone. The shift from Stage 1 to Stage 3 is better described as a shift from passive delegation to active engagement across a wider range of interaction types. Third, user population diversity: Stage 1 use is concentrated among student populations and early technical adopters; Stage 2 use is concentrated among professionally skilled workers in specific sectors; Stage 3 use has diffused to general consumers, producing a user base diverse enough in its purposes and backgrounds that no single use case dominates.

The stages are not discrete categories with hard boundaries. They are clusters within a continuous two-dimensional space, and individual countries sit at varying distances from the center of each cluster. Several countries span two stages, and their positioning is discussed explicitly. The three-stage typology maps onto the AI Matrix's quadrant logic as follows: education-first adoption is consistent with Passive Dependency; professional adoption is consistent with Elite Empowerment; and consumer diffusion is consistent with movement toward Full Empowerment. The correspondence is theoretical rather than explicitly definitional.

Before presenting the stages, one data point requires exclusion. Seychelles shows 92.3% work usage against a conversation volume of approximately 24,700, figures that are almost certainly an artifact of misclassification rather than genuine user behavior. Its task success rate of 49.2% and education level of 16.97 years, both well outside any plausible national profile, confirm that the data for this entry should not be treated as representative. It is excluded from the stage analysis.

**Figure 3. Interaction Profile by Adoption Stage**  
**Selected countries, AEI November 2025 (n > 1,000 conversations)**



Note. Four metrics are shown: work use, coursework use, directive interaction, and task iteration, each as a share of total conversations. Stage groupings are indicated by the shaded bands below the country labels. Countries are ordered within each stage by sample size. Full data for all 117 countries with  $n \geq 200$  and complete variable coverage is provided in Appendix Table A1.

Figure 3 presents the interaction profiles that define the three stages. The most diagnostic contrasts are between the coursework bars (dominant in Stage 1, minimal in Stage 3) and the directive versus iteration bars (high directive and low iteration in Stages 1 and 2, converging toward parity in Stage 3). South Korea's anomalously high iteration bar relative to its Stage 2 peers is visible here and is discussed in Section 5.3 as a leading indicator of transition toward Stage 3. Readers should note that the y-axis captures share of conversations, not absolute volume. The large US sample means its Stage 3 bars represent far more interactions than the corresponding bars for Sweden or Norway.

## 5.2 Stage 1: Education-First Adoption (Passive Dependency)

The clearest Stage 1 signal in the data is a high proportion of coursework interactions combined with elevated directive rates, below-average prompt education levels, and a use case distribution in which work-related tasks account for a minority of conversations. Countries in this cluster have reached Claude, typically through student populations and early technical adopters, before AI use has diffused into professional or personal domains at scale. This pattern is consistent with the AI Matrix's Passive Dependency quadrant. Access is adequate for the user population that has found its way to the platform, but engagement remains shallow and the iterative, professional use that characterizes higher-agency adoption has not yet developed.

Indonesia is the most prominent Stage 1 country in the dataset and the clearest illustration of the Passive Dependency dynamic. With 58.8% of conversations classified as coursework and only 22.2% as work-related, it is the most education-dominated large user base in the data. Its directive rate of 39.8% is among the highest of any country with more than 10,000 conversations, and its prompt education level of 12.2 years, while at the global mean, reflects a user base concentrated among students rather than distributed across the professional and personal range seen in Stage 3 countries. The conversation volume of over 25,000 establishes that this is a robust pattern rather than a small-sample artifact.

Latin America provides the broadest regional illustration of Stage 1. Peru shows 47.8% coursework, with only 32.8% work use and a directive rate of 36.0%. Ecuador shows 54.4% coursework. Bolivia shows 50.8%. Colombia, with 36.3% coursework and a 34.3% directive rate, sits at the upper boundary of Stage 1, with its work use at 40.5% beginning to signal transition toward Stage 2. The consistency of this pattern across a geographically dispersed set of middle-income countries supports the interpretation that the education-first cluster reflects a structural feature of adoption trajectories rather than idiosyncratic national circumstances. The directive rates across these countries range from 34 to 40%, compared to a global mean of 31.7%, indicating that when users in Stage 1 countries do engage with Claude, they are more likely to delegate a task and wait for an output than to refine, iterate, or learn.

Sub-Saharan Africa shows Stage 1 characteristics across the countries represented with sufficient sample sizes. Cameroon shows 39.5% coursework and a directive rate of 41.7%, the highest of any country in this cluster. Ivory Coast shows 36.2% coursework and a directive rate of 43.1%. Rwanda, with 35.2% coursework and a work share of 42.1%, occupies a transitional position similar to Colombia, with early signs of professional adoption emerging alongside a still-dominant education-first base. Kenya shows 32.1% coursework, and with 42.2% work and a task success rate of 63.7%, the lowest in the African subsample, it illustrates a pattern in which work use is growing but has not yet reached the interaction sophistication characteristic of Stage 2.

Three features of the Stage 1 data deserve comment beyond the headline coursework figures. First, task success rates in Stage 1 countries are not uniformly low. Indonesia's rate of 70.4% and Peru's of 71.6% are above the global mean of 66.9%. This initially appears inconsistent with a low-agency characterization, but the explanation lies in task selection: coursework tasks, particularly those involving straightforward explanation, summarization, or worked examples, tend to generate higher success rates precisely because they are structurally simpler than the multi-step professional tasks that dominate Stage 2 and Stage 3 usage. Second, the autonomy scores of Stage 1 countries are not dramatically different from the global mean. Third, Stage 1 countries do

show meaningful task iteration rates in the range of 25 to 33%, indicating that iterative engagement is not absent but is not the dominant mode. The pattern is one of degree rather than categorical difference.

### **5.3 Stage 2: Professional Adoption (Elite Empowerment)**

The Stage 2 cluster is defined by work dominance, above-average directive rates relative to income level, and a user base concentrated among technically skilled populations. Coursework is present but at lower levels than Stage 1. Personal use is limited relative to Stage 3, indicating that AI has diffused into the professional sphere without yet becoming a general consumer tool. This pattern is consistent with the AI Matrix's Elite Empowerment quadrant: a professional minority is engaging with AI intensively, at sophistication levels that generate real productivity gains, while the broader population remains below the agency threshold for meaningful participation.

India is the defining Stage 2 case and the most analytically important country in the entire dataset for the Elite Empowerment argument. With 58,098 conversations it is the second largest national sample, representing approximately 5.8% of the global total. Its work share of 51.3% and autonomy score of 3.6, the highest of any major country, are consistent with a professional user base engaging with demanding tasks at a high level of delegation. Yet several indicators qualify this apparent sophistication. The directive rate of 35.2% is higher than any Stage 3 country with comparable volume, the personal use share of 27.8% is well below the Stage 3 range, and the task success rate of 63.6% is the lowest of any country with more than 50,000 conversations. Most tellingly, Anthropic's India country brief establishes that per-capita usage ranks 101st out of 116 countries with sufficient sample, and that more than half of all Indian interactions are concentrated in four IT-sector states (Appel et al., 2026b). The high aggregate volume is a function of population scale and IT sector concentration, not broad national diffusion. This is Elite Empowerment in its most legible form. The distributional failure is visible in the data at every level of analysis.

Brazil shows the highest work share of any major country in the Stage 2 cluster at 55.2%, with a directive rate of 35.6% and a personal use share of 33.0%, the latter indicating that consumer diffusion has begun but has not yet achieved the balance characteristic of Stage 3. Its task success rate of 66.9% is at the global mean. Pakistan presents a more constrained version of the Stage 2 profile: 49.4% work, a directive rate of 40.5% that is the highest of any Stage 2 country with substantial volume, a task success rate of 60.8% that is the lowest in the cluster, and a prompt education level of 11.48 years below the global mean. This combination of high work share, high delegation, and below-average success and education levels, is consistent with intensive professional use that has not yet developed the interaction sophistication associated with higher productivity outcomes.

Vietnam occupies a transitional position between Stage 1 and Stage 2 that warrants specific discussion. Its work share of 48.0% places it at the Stage 2 threshold, but its coursework share of 24.4% is high for a country with this level of work dominance. The directive rate of 32.5% is at the lower end of Stage 2, and the iteration rate of 27.9% is above the Stage 2 mean. OpenAI's data identifies Vietnam as a high user of agentic tools relative to its income level (OpenAI, 2026b), which the AEI data partially corroborates in the work and autonomy figures, though the coursework share and modest iteration rate suggest the country is in earlier transition than the OpenAI finding alone would imply.

South Korea represents the most analytically interesting boundary case in the dataset. Its work share of 51.1% and limited personal use of 33.3% place it squarely in Stage 2, consistent with a professional adoption profile. But its task iteration rate of 35.1% is the highest of any large country in the entire dataset, well above both the Stage 2 mean and the Stage 3 range for comparable countries. Its directive rate of 28.0% is lower than the global mean and comparable to Stage 3 countries. This combination, work concentration with high iteration and low directive rates, is consistent with a country in active transition from Elite Empowerment toward Full Empowerment, where the professional user base has developed sophisticated interaction patterns that have not yet diffused to produce the personal use balance characteristic of Stage 3. The per-capita analysis reinforces this reading: South Korea's 592 conversations per million population places it among the highest-reach countries in the dataset, comparable to Australia and Canada, suggesting that access has become broadly distributed even before interaction patterns have fully shifted. South Korea's trajectory warrants monitoring in future AEI waves as a potential leading indicator of how Elite Empowerment transitions toward Full Empowerment.

#### **5.4 Stage 3: Consumer Diffusion (Movement Toward Full Empowerment)**

The Stage 3 cluster is defined by a balanced use case distribution in which personal use constitutes a substantial share alongside work, lower directive rates, and a distributional spread of users broad enough to encompass general consumers rather than primarily students and professionals. Countries in this cluster have moved beyond the point where AI is primarily a tool for specific occupational or educational purposes. They show interaction patterns consistent with higher agency, particularly lower directive rates and higher task iteration, and their user populations are more diverse in terms of the tasks they bring to the platform. This pattern is consistent with movement toward the AI Matrix's Full Empowerment quadrant, though the data cannot establish whether any country has achieved that position rather than approaching it.

The Scandinavian countries collectively provide the clearest Stage 3 signal in the data. Norway shows 43.8% personal use, a directive rate of 25.0%, and a task success rate of 68.1%. Sweden shows 42.8% personal use and a directive rate of 22.9%, the lowest of any country with more than 1,000 conversations in the dataset. Finland shows 43.1% personal use and a directive rate of 24.7%. Denmark shows 42.6% personal use and a directive rate of 23.5%. These four countries form a coherent cluster at the lower end of the directive range and the upper end of personal use, consistent with a population in which AI has diffused beyond the professional and educational contexts that dominate earlier stages.

Germany is notable for having the highest personal use of any major country in the dataset at 44.7%, which exceeds its work share of 41.5%, making it the clearest case among large economies of personal use having overtaken professional use. Its directive rate of 26.6% matches the Stage 3 pattern. The UK's directive rate of 27.4% is the lowest of any country with more than 30,000 conversations, and its iteration rate of 30.5% is among the highest of Stage 3 countries, consistent with a user base that has developed the iterative engagement habits the framework associates with high agency. Australia matches the UK pattern closely at 28.2% directive and 30.5% iteration across a sample of over 15,000 conversations.

The United States presents a nuanced Stage 3 profile. With 219,666 conversations, it is by far the largest national sample, accounting for approximately 22% of all observations. Its personal use share of 40.4% and work share of 41.9% are closely balanced, consistent with broad consumer diffusion. However, its directive rate of 31.9% is higher than most other Stage 3 countries, and its iteration rate of 26.7% is at the lower end of the Stage 3 range. Its autonomy score of 3.25 is the lowest of any major country. This combination is consistent with a very large and diverse user population in which the sophistication range is wider than in smaller Stage 3 countries, pulling aggregate interaction patterns toward the mean. The within-country convergence finding, which Anthropic estimates would equalize per-capita usage across US states within two to five years (Appel et al., 2026a), is consistent with a population broadly above the agency threshold whose heterogeneity of engagement reflects distributional breadth rather than concentrated elite use.

Singapore sits at a distinct position across all three stages. Its personal use of 33.0% and coursework share of 26.2% give it a mixed Stage 2 to Stage 3 profile, but its prompt education level of 12.83 years is the highest of any country in the key comparators, well above the global mean of 12.2 years, and its directive rate of 28.2% is at the lower end of Stage 3. On the per-capita measure, Singapore leads the entire dataset at 863 conversations per million population, considerably above Israel (817), the United States (656), and Canada (597). Singapore's position in the AI Matrix is best understood as a country where access is not merely present but deeply

embedded at the population level, agency is high among the educated professional population, and consumer diffusion is still developing toward the personal use balance seen in the Nordic countries. The socio-economic resilience argument in Section 3.2 is consistent with this profile. Both Figures 2a and 2b in Section 4.3 make Singapore's position visible directly: it leads the per-capita distribution while sitting in the upper-right zone of Figure 2a by virtue of absolute volume.

## **5.5 The Within/Between Asymmetry**

The most theoretically important finding in the AEI data is not visible in any individual country's figures. It lies in the relationship between within-country convergence and between-country divergence. Within the United States, Anthropic estimates that per-capita AI usage would equalize across states within two to five years, a diffusion speed roughly ten times faster than comparable technologies in the twentieth century (Appel et al., 2026a). The data in this section shows that the global distribution of countries across the three adoption stages shows no equivalent convergence. The distance between Indonesia's 58.8% coursework and Germany's 44.7% personal use is not narrowing in the cross-sectional data available here. The distance between Pakistan's 40.5% directive rate and Sweden's 22.9% is not a gap that infrastructure provision alone will close.

The AI Matrix provides the theoretical account of this asymmetry that the capability overhang literature does not offer. Within the United States, all states share a governance framework, a labor market structure, a regulatory environment, and a language. Once the access threshold is met, diffusion is driven by workforce composition and institutional learning, both of which operate rapidly within these shared conditions. Between countries, the access threshold itself is a governance product that varies enormously, and the institutional conditions for agency-building, including education system quality, labor market structures that reward higher-order skills, and regulatory environments that support iterative professional AI use, vary at least as much. The capability overhang between countries is not narrowing because the conditions that would narrow it are not converging. They are, in several respects, compounding.

The three-stage typology developed in this section is a snapshot of this compounding process. Education-first countries are building familiarity with AI through student populations, but without the professional and institutional infrastructure to convert that familiarity into the iterative, high-agency engagement that generates compounding productivity returns. Professional adoption countries have developed that infrastructure within a technical elite but have not yet created the distributional conditions for it to diffuse. Consumer diffusion countries have moved beyond both constraints, but they represent a minority of the 173 countries in the data and are disproportionately concentrated in high-income regions. The policy implications of this distribution are taken up in Section 6.

## 6. Discussion

### 6.1 Policy Implications of a Two-Dimensional Framework

The three-stage typology in Section 5 is not merely a descriptive exercise. It carries a direct and uncomfortable implication for the policy consensus that has formed around the capability overhang problem. If access and agency are genuinely distinct dimensions, each capable of failing independently and each requiring a different intervention, then the standard policy response, more infrastructure, more AI skills training, more institutional readiness, will produce different outcomes depending on where a country sits in the two-dimensional space. In some contexts these interventions are precisely what is needed. In others they will expand access to tools that remain underused, or build technical literacy among populations who still cannot reach the tools. Getting the diagnosis wrong is not a neutral error. It allocates scarce resources to interventions that cannot close the gap they are designed to address, while the compounding productivity advantages of countries already in Stage 3 continue to accumulate.

For countries in Stage 1, the education-first pattern is a signal that AI has arrived through the path of least resistance: students and early technical adopters who have found their way to the platform independently. For the user populations currently present in the data, the interaction patterns suggest that the platform is reachable: the Stage 1 signal is not that users cannot get to Claude, but that when they do, they engage with it directly and narrowly. This does not establish that national-level access constraints have ceased to be primary for the broader population not yet represented in the data. It does suggest that for the slice of the population already above the access threshold, the more pressing constraint is the absence of institutional and professional conditions that would convert student familiarity into iterative, work-oriented engagement. For that population, the appropriate intervention is agency-building: integration of AI into professional development frameworks, curriculum redesign that treats prompting and critical evaluation of AI outputs as foundational skills rather than supplementary ones, and labor market policies that create demand for higher-order AI engagement beyond the student population. Investing primarily in further access expansion for these countries is likely to move more users to the platform without changing what they do there. But this inference applies within the observed frontier. It does not license a settled claim that infrastructure is no longer the primary national bottleneck in Stage 1 countries, many of whose populations remain absent from the data entirely.

For countries in Stage 2, the distributional failure is the central problem. The professional adoption pattern in India, Brazil, Pakistan, and Thailand reflects real, sophisticated AI use by an identifiable technical elite. The productivity gains that this use generates are substantial and well-documented. But they are accruing to a narrow occupational and geographic base, and the data

shows no evidence that diffusion to the broader population is occurring spontaneously. The appropriate intervention is structural broadening: deliberate extension of AI literacy and professional development to sectors and regions currently outside the IT and finance concentrations that dominate Stage 2 usage, economic policies that create AI-adjacent employment opportunities in healthcare, education, agriculture, and public administration, and governance frameworks that ensure the benefits of AI adoption are not captured exclusively by the existing technical elite. Additional access investment in Stage 2 countries may be appropriate at the margins, particularly for the non-IT populations in rural or lower-income regions, but the primary constraint is distributional rather than infrastructural.

For countries in Stage 3, the policy challenge is one of maintenance and governance rather than transition. The within-country convergence finding for the United States suggests that once a population is broadly above the agency threshold and the institutional conditions for diffusion are in place, AI adoption proceeds rapidly through workforce and institutional channels. The risk in Stage 3 countries is regression rather than stagnation: the erosion of AI literacy through passive consumption habits, the concentration of AI capability in large platforms without adequate regulatory oversight, or the widening of within-country distributional gaps as AI-intensive industries pull further ahead of AI-excluded ones. These are governance problems rather than access or agency problems, and they require regulatory and institutional responses rather than additional investment in access or skills.

Against this differentiated picture, the OpenAI for Countries initiative, which proposes education partnerships, health sector AI deployment, AI skills training, and startup accelerator programs as a standard package for participating governments (OpenAI, 2026b), reflects an approach that is reasonable in intent but analytically undifferentiated. A country in Stage 1 needs agency-building before additional access provision makes productive sense, and skills training that focuses on tool familiarity rather than iterative critical engagement will replicate the passive dependency pattern already visible in the data. A country in Stage 2 needs distributional broadening and governance frameworks that extend high-agency use beyond the technical sector, not startup accelerators that are likely to deepen the concentration of AI capability in the professional elite already driving usage. Treating all three stages with the same program risks replicating the first-level error that the digital divide literature documented over two decades: expanding the reach of technology without changing the conditions under which it is used, and measuring success by adoption rates rather than by the distribution of productive engagement.

## 6.2 The Path Dependency Problem

The differentiated policy framework above raises a hard question: why has the global adoption curve not steepened despite significant investment in AI access since 2022? The McKinsey Global Institute, the IMF, OpenAI, and Anthropic have all documented the pattern. Capability is advancing rapidly. Access is expanding. Usage rates are rising in absolute terms across all country groups. Yet the relative gaps between countries are stable, and in some comparisons widening, as Stage 3 countries extract compounding productivity returns while Stage 1 and Stage 2 countries capture a fraction of the available benefit. The AI Matrix's path dependency pillar provides the most coherent account of why this is happening.

Path dependency theory predicts exactly what the data shows. The institutional conditions that support high-agency AI use, education systems capable of producing iterative, critical AI engagement, labor market structures that reward higher-order knowledge work, and governance frameworks that support professional AI integration across sectors, are not products that can be created quickly through targeted investment. They are outputs of long developmental trajectories in which prior investment in education, institutions, and economic diversification compounds over time. Countries that enter the AI adoption period with these foundations in place, the Scandinavian cluster, Germany, the UK, Australia, Japan, can move rapidly toward Full Empowerment because the conditions for agency-building already exist at scale. Countries that enter without these foundations, regardless of how much access investment they receive, face a structural barrier that access provision cannot dissolve.

This is not a counsel of despair, and the AI Matrix is explicit that path dependencies can be broken. Singapore's trajectory is the most instructive counter-example available. Starting from a position in the 1990s where the institutional conditions for high-agency technology use were limited, Singapore pursued a sustained, coordinated program of investment across the education-governance-infrastructure triad that compressed a developmental path that typically takes decades into approximately two decades (Smart Nation and Digital Government Office, 2023). The AEI data reflects the result: a prompt education level of 12.83 years, the highest of any country in the key comparators, a directive rate of 28.2% consistent with Stage 3 rather than Stage 2, and at 863 conversations per million population, the highest per-capita reach in the entire dataset.

What Singapore illustrates, however, is also what makes the path dependency problem so resistant to standard policy tools. The investment that produced Singapore's current AI adoption profile was not a technology policy. It was a comprehensive institutional development strategy involving education reform, labor market restructuring, governance capacity-building, and

sustained public investment over a multi-decade period. Replicating this trajectory requires a form of coordinated, sustained commitment that most international AI policy frameworks, including the OpenAI for Countries initiative and the multilateral commitments made at various international AI summits in 2025, do not currently contemplate at the required scale or over the required time horizon. The global capability overhang is not narrowing because narrowing it requires not just technology policy but the broader institutional development that makes technology productive, and that development cannot be accelerated primarily through access provision.

### **6.3 Limitations and Future Research**

The empirical analysis in this paper rests on a single provider's data covering one week of interactions, and these constraints set a clear ceiling on the strength of the claims the analysis can support. The AEI dataset captures only countries already above the access threshold. The countries most likely to be in Full Dependency, and therefore most in need of the framework's policy prescriptions, are largely absent from the data or present in sample sizes too small for reliable analysis. The empirical contribution is within-frontier variation: a characterization of how countries that have reached Claude differ from each other in patterns of engagement, not a complete map of global AI adoption.

The agency composite is a behavioral proxy, constructed from observable interaction patterns, rather than a validated measure of agency as the AI Matrix conceptualizes it. It is possible that the collaboration pattern classifications produced by Clio's automated analysis reflect platform design features or task type distributions rather than underlying differences in user sophistication. The cross-sectional design cannot distinguish between countries that are stable in their current stage and countries in active transition.

A brief robustness note on the stage assignments is warranted. The three stage thresholds, coursework share above 30% for Stage 1, work share above 48% with coursework below 25% for Stage 2, and personal use above 38% for Stage 3, were set to reflect theoretical distinctions rather than optimized statistical cutpoints. Modest variation in these thresholds does not materially alter the country groupings or the broad distributional pattern. Shifting the Stage 1 coursework threshold between 25% and 35% reassigns fewer than eight countries, all of which are genuine boundary cases discussed explicitly in Section 5. The core clusters, Indonesia-Peru-Ecuador in Stage 1, India-Brazil-Pakistan in Stage 2, and the Scandinavian-German-UK-US group in Stage 3, are robust to threshold variation across the plausible range. The residual classification rule, which assigns countries not meeting any primary threshold by agency score, handles approximately 20%

of the reliable sample; alternative thresholds for the agency residual cutpoints produce equivalent reassignment rates to those seen in the primary rule variation. Longitudinal analysis across multiple AEI releases, examining whether the stage assignments proposed here are stable over time and whether the within/between asymmetry in convergence speed is widening or narrowing, is an important direction for future work that extends beyond the scope of the present paper. A systematic cross-provider comparison using matched country samples and standardized variable definitions would be a considerably more powerful test of whether the three-stage typology is a feature of AI adoption in general or a product of Claude's specific user base. The Brynjolfsson et al. (2025) working paper on employment effects, which uses earlier AEI data, demonstrates that the dataset is already generating independent external research. As additional waves accumulate and other providers release comparable interaction data, the empirical foundation for a framework of this kind will become substantially more robust.

## **7. Conclusion**

The capability overhang is real and consequential. The evidence from Anthropic, OpenAI, and Microsoft points consistently toward the same structural finding. The gap between what AI systems can do and what most people actually do with them is large, geographically structured, and showing no signs of spontaneous closure. Whether AI capability is now sufficient for a wide range of knowledge work tasks is a question reasonable observers still debate. What is not seriously in dispute is that the depth and character of AI engagement varies enormously across countries in ways that are not explained by access alone.

What this paper has argued is that the capability overhang concept, as currently deployed, is not analytically tractable enough to guide the policy response this problem requires. By treating the gap between AI's potential and its realization as a single undifferentiated phenomenon, it generates undifferentiated prescriptions. More access, more education and more infrastructure are reasonable starting points, but they are not differentiated to the specific situation of countries at different positions on the adoption curve, and applying them without that differentiation risks replicating, in AI, the first-level error that the digital divide literature documented over two decades of internet policy. Expanding access to tools that remain underused does not close the gap. In some configurations it widens it, by delivering capability to populations without the institutional conditions to use it productively while Stage 3 countries continue to extract compounding returns from increasingly sophisticated engagement.

The first step in getting that intervention right is knowing which stage you are in. The AI Matrix, operationalized against the most comprehensive observed dataset of AI interaction behavior currently available, provides one tool for making that determination. It is not the only tool, and the limitations documented in Section 4 establish clearly what a stronger empirical foundation would require. But the convergence of the AEI data and OpenAI's independently collected evidence, interpreted through a framework that preceded both, is a meaningful form of evidence. The capability overhang will not be closed by access provision alone. It will be closed by the patient, differentiated, institutionally grounded work of building the conditions under which AI becomes not merely usable but genuinely useful, to more people, in more places, in ways that compound over time rather than concentrate.

## References

- Acemoglu, D., Johnson, S., & Robinson, J. A. (2005). Institutions as a fundamental cause of long-run growth. In P. Aghion & S. N. Durlauf (Eds.), *Handbook of Economic Growth* (Vol. 1A, pp. 385-472). Elsevier. [https://doi.org/10.1016/S1574-0684\(05\)01006-3](https://doi.org/10.1016/S1574-0684(05)01006-3)
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3-30. <https://doi.org/10.1257/jep.33.2.3>
- Appel, R., Aldous, S., Eaton, J., Ho, H., Massenkoff, M., McCrory, P., Mueller, J., Pastewka, E., Pollack, S., Rajkumar, N., Saunders, D., Sanderford, A., & Withee, K. (2025). *Anthropic Economic Index report: Uneven geographic and enterprise AI adoption*. Anthropic. <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>
- Appel, R., Aldous, S., Eaton, J., Ho, H., Massenkoff, M., McCrory, P., Mueller, J., Pastewka, E., Pollack, S., Rajkumar, N., Saunders, D., Sanderford, A., & Withee, K. (2026a). *Anthropic Economic Index report: Economic primitives*. Anthropic. <https://www.anthropic.com/research/anthropic-economic-index-january-2026-report>
- Appel, R., Aldous, S., Eaton, J., Ho, H., Massenkoff, M., McCrory, P., Mueller, J., Pastewka, E., Pollack, S., Rajkumar, N., Saunders, D., Sanderford, A., & Withee, K. (2026b). *India country brief: The Anthropic Economic Index*. Anthropic. <https://www.anthropic.com/research/india-brief-economic-index>
- Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394), 116-131. <https://doi.org/10.2307/2234208>
- Brynjolfsson, E., Chandar, B., & Chen, R. (2025). *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*. Stanford Digital Economy Lab Working Paper. <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/>

- David, P. A. (1985). Clio and the economics of QWERTY. *American Economic Review*, 75(2), 332-337.
- Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
- Feenberg, A. (1991). *Critical theory of technology*. Oxford University Press.
- Folke, C., Carpenter, S. R., Walker, B., Scheffer, M., Chapin, T., & Rockström, J. (2010). Resilience thinking: Integrating resilience, adaptability and transformability. *Ecology and Society*, 15(4), Article 20.
- Hargittai, E. (2002). Second-level digital divide: Differences in people's online skills. *First Monday*, 7(4). <https://doi.org/10.5210/fm.v7i4.942>
- Holling, C. S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4, 1-23. <https://doi.org/10.1146/annurev.es.04.110173.000245>
- International Monetary Fund. (2024). *Gen-AI: Artificial intelligence and the future of work*. IMF Staff Discussion Note SDN/2024/001.
- Jessop, B. (2002). Liberalism, neo-liberalism, and urban governance: A state-theoretical perspective. *Antipode*, 34(3), 452-472.
- Microsoft AI Economy Institute. (2026). *Global AI adoption in 2025: A widening digital divide*. Microsoft. <https://aka.ms/msft-ai-diffusion-report-2025-h2>
- Norris, P. (2001). *Digital divide: Civic engagement, information poverty, and the internet worldwide*. Cambridge University Press.
- OpenAI. (2026a). *Ending the capability overhang*. OpenAI. <https://cdn.openai.com/pdf/openai-ending-the-capability-overhang.pdf>
- OpenAI. (2026b). *How countries can end the capability overhang*. OpenAI. <https://openai.com/index/how-countries-can-end-the-capability-overhang/>
- Ostrom, E. (2009). A polycentric approach for coping with climate change. *World Bank Policy Research Working Paper* 5095.
- Oxford Insights. (2024). *Government AI Readiness Index 2024*. Oxford Insights. <https://www.oxfordinsights.com/government-ai-readiness-index-2024>
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.
- Selwyn, N. (2004). Reconsidering political and popular understandings of the digital divide. *New Media & Society*, 6(3), 341-362. <https://doi.org/10.1177/1461444804042519>

- Simpson, E. (2025a). The AI Matrix: Empowerment or dependency? A conceptual framework. SSRN. <https://doi.org/10.2139/ssrn.5228571>
- Smart Nation and Digital Government Office. (2023). Transforming Singapore through AI and digital innovation. <https://www.smartnation.gov.sg>
- Stanford Institute for Human-Centered Artificial Intelligence. (2024). AI Index Report 2024. Stanford University. <https://aiindex.stanford.edu/report/2024/>
- Stoker, G. (1998). Governance as theory: Five propositions. *International Social Science Journal*, 50(155), 17-28.
- Tortoise Media. (2023). The Global AI Index 2023. Tortoise Media. <https://www.tortoisemedia.com/intelligence/global-ai/>
- van Deursen, A. J. A. M., & van Dijk, J. A. G. M. (2014). The digital divide shifts to differences in usage. *New Media & Society*, 16(3), 507-526. <https://doi.org/10.1177/1461444813487959>
- van Dijk, J. A. G. M. (2005). *The deepening divide: Inequality in the information society*. SAGE Publications.
- Warschauer, M. (2003). *Technology and social inclusion: Rethinking the digital divide*. MIT Press.
- World Bank. (2024). World development indicators: Population, total. World Bank Open Data. <https://data.worldbank.org/indicator/SP.POP.TOTL>

## **Appendix. Variable Construction and Country Data**

This appendix has two parts. The first explains the derivation of the two composite indicators used in Figure 2 and in the three-stage typology: the Human Agency Index and the Access to AI Resources index. The second presents the full country-level data for all 117 countries with  $n \geq 200$  conversations and complete variable coverage in the AEI November 2025 dataset, together with each country's computed composite scores and stage assignment.

### **A1. Derivation of Composite Indicators**

#### **Human Agency Index**

The Human Agency Index is a country-level composite constructed from four observable interaction variables drawn from the AEI dataset. Each variable functions as a behavioral proxy for the capacity to engage with AI tools in iterative, critical, and productive ways, as the AI Matrix conceptualizes agency. The four components are:

(1) Task iteration rate (%). The proportion of conversations classified by Clio as task iteration: exchanges in which the user develops and refines a task across multiple turns. Higher rates indicate more agentic engagement. This is the primary behavioral signal for iterative rather than passive use.

(2) Directive rate (%), inverted. The proportion of conversations classified as directive: single-turn delegation in which the user issues an instruction and accepts the output without revision. This variable is inverted before inclusion, so that high directive rates contribute negatively to the composite. Countries where delegation dominates score lower on agency.

(3) Task success rate (%). Claude's self-assessment of whether it completed the requested task. Higher success rates on work-oriented tasks are associated with users who specify requests precisely, evaluate outputs critically, and iterate when results fall short. At the country level, this functions as a population-level outcome proxy for interaction quality.

(4) Prompt education level (years). Clio's estimate of the years of formal schooling required to write the user's prompt, on the US scale where 12 years = high school diploma and 16 years = bachelor's degree. More educationally demanding prompts tend to reflect more precise task

specification and more complex underlying needs. This variable captures a dimension of sophistication not captured by the collaboration pattern variables alone.

Construction: Each component is independently min-max normalised across all 117 countries in the reliable sample ( $n \geq 200$ , Seychelles excluded) to the  $[0, 1]$  interval. The directive rate is then inverted (1 minus the normalised value) so that all four components point in the same direction: higher values indicate higher agency. The composite is the unweighted mean of the four normalised components:

$$\text{Agency} = (N_{\text{iteration}} + (1 - N_{\text{directive}}) + N_{\text{success}} + N_{\text{education}}) / 4$$

where  $N$  denotes the min-max normalised value of each variable. The composite is ordinal: it supports comparisons of relative position across countries but should not be treated as a cardinal scale. A country scoring 0.70 is more agentive than one scoring 0.50 in the relative sense captured by these four proxies; it does not mean that 40 percent more agency is present in some absolute sense.

### **Access to AI Resources Index**

The Access to AI Resources index is a single-variable proxy derived from the log base-10 of each country's conversation volume in the AEI November 2025 dataset. Log transformation is applied to compress the extreme right-skew of the raw distribution, in which the United States accounts for approximately 22 percent of all conversations. The transformed values are then min-max normalised to the  $[0, 1]$  interval across the 117 reliable countries in the sample.

This is an imperfect measure of access in the AI Matrix sense, which refers to the structural availability of AI tools rather than observed use. Observed conversation volume is a function of both access (the probability that a person can reach Claude) and adoption (the probability that a person who can reach Claude chooses to use it). These cannot be cleanly separated within the AEI data alone. The measure is nonetheless informative as a relative indicator of effective reach or realized access in practice: countries with higher normalised scores have demonstrably larger user populations relative to the dataset, which is consistent with greater effective access. Throughout the paper, "access" on the horizontal axis should be read in this sense — as observed engagement volume rather than structural availability of infrastructure. Countries with very low scores or absent

from the dataset altogether are plausibly below an access threshold rather than simply low in propensity to adopt.

A per-capita complement to this raw volume measure is developed in Section 4.3a and presented as Figure 2b. That measure uses World Bank and UN 2024 population estimates, matched successfully for 171 of the 172 countries in the dataset, to construct conversations per million population as an alternative access proxy. Both measures are available in Table A1. The Oxford Insights Government AI Readiness Index (Oxford Insights, 2024), which would provide a structurally independent access indicator, was not available in machine-readable form at country level with sufficient coverage for direct incorporation. Future work using an independent infrastructure index alongside the AEI behavioral data would provide a stronger test of the two-dimensional framework.

### **Stage Assignment**

Countries are assigned to the three adoption stages using a sequential rule applied to the raw AEI variables rather than the composite scores, so that the stage classification is directly interpretable in terms of observable interaction patterns. The rules are applied in priority order:

Stage 1 (Education-first, Passive Dependency): coursework share > 30%.

Stage 2 (Professional adoption, Elite Empowerment): work share > 48% and coursework share < 25%.

Stage 3 (Consumer diffusion, toward Full Empowerment): personal use share > 38%.

Residual: countries not assigned by the above rules are classified by agency composite score:  $\geq 0.55$  = Stage 3;  $\geq 0.42$  = Stage 2;  $< 0.42$  = Stage 1.

These thresholds were set to reflect the qualitative boundaries described in Section 5. They should be understood as operationalizations of theoretical distinctions rather than optimised statistical cutpoints. Modest variation across the plausible range does not materially alter the broad stage structure or the core country groupings; the detailed sensitivity analysis is provided in Section 6.3.

### **Table A1. Country-Level Data and Composite Scores**

All 117 countries with  $n \geq 200$  conversations and complete variable coverage; AEI November 2025. Sorted by stage, then alphabetically. Seychelles ( $n = 24,715$ ; 92.3% work; edu = 16.97) excluded as probable misclassification artifact. An additional 54 countries appear in the dataset

with  $n < 200$  and are plotted as low-confidence observations in Figure 2; they are not included here because agency scores cannot be reliably computed at that sample size. Approximately 23 countries are absent from the dataset entirely and are annotated in the Full Dependency quadrant of Figure 2.

Country	n	Work %	Personal %	Coursework %	Directive %	Iteration %	Success %	Edu (yrs)	Agency (0-1)	Access (0-1)	Stage
Algeria	2,710	29.8	29.6	40.6	34.0	25.6	69.3	11.79	0.523	0.356	Stage 1
Angola	265	43.0	34.3	22.6	44.5	25.7	62.6	10.23	0.206	0.015	Stage 1
BF	344	35.8	23.8	40.4	33.4	25.9	64.2	11.97	0.473	0.053	Stage 1
Bahrain	281	43.4	31.0	25.6	36.3	26.0	62.6	11.65	0.397	0.024	Stage 1
Bolivia	1,536	28.3	20.8	50.8	34.6	30.5	70.1	11.81	0.612	0.273	Stage 1
Cameroon	1,028	35.7	24.8	39.5	41.7	23.2	65.8	11.91	0.353	0.214	Stage 1
Colombia	8,841	40.5	23.2	36.3	34.3	31.6	69.6	11.98	0.639	0.529	Stage 1
Costa Rica	1,073	38.8	29.5	31.7	31.0	29.3	69.0	11.90	0.622	0.220	Stage 1
Dominican Republic	1,495	33.7	28.6	37.7	39.2	27.2	68.6	10.75	0.413	0.269	Stage 1
Ecuador	3,452	26.5	19.1	54.4	36.6	33.3	74.1	11.76	0.693	0.391	Stage 1
El Salvador	536	39.4	26.1	34.5	29.9	28.7	72.2	11.79	0.664	0.118	Stage 1
Honduras	393	35.9	30.8	33.3	33.6	28.0	72.5	11.36	0.586	0.073	Stage 1
Indonesia	25,075	22.2	19.0	58.8	39.8	32.4	70.4	12.19	0.616	0.682	Stage 1
Iraq	3,234	13.9	47.8	38.2	34.2	20.4	77.7	9.06	0.373	0.382	Stage 1
Ivory Coast	1,141	37.7	26.1	36.2	43.1	24.6	67.4	11.67	0.368	0.229	Stage 1
Jamaica	314	35.7	33.1	31.2	38.9	31.2	71.0	11.73	0.584	0.040	Stage 1
Kenya	3,940	42.2	25.7	32.1	37.0	28.8	63.7	12.53	0.511	0.411	Stage 1
Laos	272	47.1	32.4	20.6	41.5	23.9	72.8	10.46	0.373	0.019	Stage 1
Mexico	10,345	34.5	32.8	32.6	33.1	29.7	69.7	11.43	0.585	0.552	Stage 1
Mozambique	366	32.5	32.2	35.2	41.3	27.9	65.6	10.38	0.333	0.062	Stage 1
Panama	760	34.5	31.6	33.9	36.6	29.9	71.1	11.17	0.552	0.170	Stage 1
Peru	7,375	32.8	19.4	47.8	36.0	32.8	71.6	11.88	0.662	0.503	Stage 1
Philippines	9,935	40.2	27.3	32.5	34.7	32.9	69.1	11.74	0.633	0.546	Stage 1
Rwanda	677	42.1	22.6	35.2	39.3	27.5	64.1	12.04	0.437	0.153	Stage 1
Senegal	787	43.1	25.8	31.1	34.3	27.7	68.0	11.72	0.532	0.175	Stage 1
Sri Lanka	3,190	44.6	24.6	30.8	40.2	23.4	65.4	12.00	0.373	0.380	Stage 1
TG	248	37.9	29.0	33.1	45.2	22.6	64.1	11.45	0.248	0.005	Stage 1
Tunisia	2,500	39.2	23.0	37.8	31.1	27.9	65.2	12.43	0.577	0.344	Stage 1
Uganda	661	42.4	25.4	32.2	37.4	27.2	67.6	12.16	0.512	0.149	Stage 1
Ukraine	7,489	36.0	32.2	31.8	34.6	27.4	70.2	11.62	0.549	0.505	Stage 1
Uzbekistan	827	28.4	34.2	37.2	34.6	27.6	69.0	10.48	0.459	0.182	Stage 1
Zambia	374	38.2	27.5	34.2	38.5	27.3	69.8	12.12	0.531	0.066	Stage 1
Zimbabwe	365	39.7	28.2	32.1	44.1	24.1	62.7	12.44	0.331	0.062	Stage 1
Albania	329	56.2	30.7	13.1	26.7	27.1	70.8	12.08	0.671	0.047	Stage 2
Azerbaijan	563	42.3	36.2	21.3	29.0	24.9	64.8	11.42	0.477	0.126	Stage 2
Bangladesh	2,914	39.7	31.4	28.9	35.5	26.1	68.4	11.49	0.482	0.367	Stage 2
Bosnia	334	55.4	32.0	12.6	29.6	22.5	71.6	11.90	0.560	0.049	Stage 2
Brazil	27,602	55.2	33.0	11.8	35.6	30.9	66.9	11.51	0.542	0.696	Stage 2

Country	n	Work %	Personal %	Coursework %	Directive %	Iteration %	Success %	Edu (yrs)	Agency (0-1)	Access (0-1)	Stage
Bulgaria	1,465	48.2	40.8	11.1	28.5	24.7	67.8	11.40	0.522	0.266	Stage 2
CY	547	54.5	37.8	7.7	32.2	28.0	63.6	11.85	0.505	0.121	Stage 2
Cambodia	571	49.0	36.3	14.7	33.8	24.0	66.2	11.19	0.413	0.128	Stage 2
Egypt	5,501	45.6	29.2	25.2	31.9	24.8	64.4	11.92	0.470	0.460	Stage 2
Ghana	1,546	40.8	30.7	28.5	37.7	27.6	66.8	11.98	0.492	0.274	Stage 2
India	58,098	51.3	27.8	20.9	35.2	27.4	63.6	12.25	0.488	0.805	Stage 2
Jordan	752	44.7	31.9	23.4	33.6	24.6	64.1	11.80	0.435	0.168	Stage 2
LU	386	53.1	36.3	10.6	31.6	28.2	63.7	12.41	0.554	0.070	Stage 2
Lebanon	776	44.6	28.2	27.1	31.6	25.6	65.2	12.07	0.509	0.173	Stage 2
MT	272	52.6	36.4	11.0	28.3	28.3	63.2	12.05	0.561	0.019	Stage 2
Madagascar	661	52.6	26.8	20.6	34.9	28.9	60.5	12.03	0.457	0.149	Stage 2
Mauritius	302	52.6	32.8	14.6	33.8	24.5	70.2	11.69	0.513	0.034	Stage 2
Moldova	606	39.8	32.8	27.4	31.4	26.6	66.5	11.65	0.519	0.136	Stage 2
Morocco	5,166	38.9	31.6	29.5	32.1	28.3	66.8	11.62	0.543	0.450	Stage 2
Nepal	1,673	50.7	28.2	21.0	33.2	28.8	69.2	12.28	0.617	0.285	Stage 2
Nigeria	7,629	47.4	35.2	17.4	37.3	30.8	68.5	11.56	0.547	0.508	Stage 2
North Macedonia	339	52.8	27.7	19.5	32.4	24.5	66.7	12.25	0.515	0.051	Stage 2
Pakistan	8,916	49.4	29.3	21.2	40.5	26.7	60.8	11.48	0.325	0.530	Stage 2
Paraguay	571	44.7	32.2	23.1	32.6	27.0	68.7	11.71	0.548	0.128	Stage 2
Poland	9,040	48.2	40.0	11.8	28.7	25.0	67.6	12.04	0.564	0.532	Stage 2
Saudi Arabia	3,136	43.4	35.3	21.4	33.1	26.2	63.2	11.66	0.446	0.377	Stage 2
Serbia	1,391	52.6	39.5	8.0	26.5	25.6	68.1	11.91	0.598	0.258	Stage 2
South Africa	4,459	48.6	34.1	17.2	35.0	27.9	65.1	12.07	0.508	0.429	Stage 2
South Korea	30,618	51.1	33.3	15.6	28.0	35.1	68.4	12.09	0.759	0.711	Stage 2
Tanzania	332	53.0	35.8	11.1	34.9	23.2	65.7	11.58	0.406	0.048	Stage 2
Thailand	8,032	48.9	33.3	17.7	34.8	28.5	68.1	11.67	0.538	0.515	Stage 2
UAE	3,883	54.5	30.8	14.6	32.4	32.2	64.6	11.95	0.595	0.409	Stage 2
Argentina	6,088	42.9	33.9	23.2	29.7	26.0	69.5	11.68	0.574	0.474	Stage 3
Armenia	533	46.0	34.9	19.1	34.0	28.0	70.2	11.59	0.564	0.118	Stage 3
Australia	15,488	43.9	41.7	14.4	28.2	30.5	67.0	12.04	0.654	0.611	Stage 3
Austria	3,036	43.7	39.5	16.9	26.4	26.8	68.2	12.10	0.633	0.373	Stage 3
BJ	536	47.8	27.8	24.4	35.6	30.8	67.5	11.89	0.574	0.118	Stage 3
Belgium	4,421	45.4	40.3	14.2	26.7	27.7	67.5	12.01	0.629	0.428	Stage 3
Canada	22,854	38.4	41.8	19.8	26.9	28.0	67.5	12.19	0.644	0.668	Stage 3
Chile	3,819	42.9	33.0	24.1	31.2	28.7	67.8	11.97	0.597	0.406	Stage 3
Croatia	861	46.5	43.1	10.5	29.5	23.3	67.6	11.74	0.506	0.188	Stage 3
Czech Republic	2,888	43.6	43.2	13.2	27.6	25.4	66.4	11.82	0.551	0.365	Stage 3
Denmark	2,151	44.2	42.6	13.1	23.5	26.1	67.0	11.91	0.624	0.322	Stage 3
Estonia	713	40.5	47.7	11.8	25.2	23.0	68.3	11.85	0.567	0.160	Stage 3
Finland	1,821	40.3	43.1	16.6	24.7	27.5	70.0	12.28	0.702	0.298	Stage 3
France	30,208	41.7	40.2	18.1	28.4	29.8	67.6	11.87	0.638	0.709	Stage 3
Georgia	1,369	44.6	36.4	19.0	27.4	27.1	68.4	11.71	0.604	0.256	Stage 3
Germany	25,382	41.5	44.7	13.8	26.6	26.6	67.6	11.78	0.598	0.684	Stage 3
Greece	2,137	46.7	37.4	15.9	28.6	26.1	65.1	12.37	0.569	0.321	Stage 3
Guatemala	808	39.7	30.8	29.5	30.8	29.2	70.8	11.78	0.641	0.179	Stage 3

Country	n	Work %	Personal %	Coursework %	Directive %	Iteration %	Success %	Edu (yrs)	Agency (0-1)	Access (0-1)	Stage
Hungary	1,643	40.4	44.3	15.2	28.8	25.3	65.2	11.66	0.508	0.283	Stage 3
IE	2,275	40.6	39.2	20.2	29.1	26.2	67.3	12.26	0.590	0.330	Stage 3
Israel	7,920	45.1	41.6	13.2	28.0	29.2	67.1	11.37	0.592	0.513	Stage 3
Italy	16,336	44.9	39.0	16.1	29.3	28.9	68.9	11.63	0.615	0.619	Stage 3
Japan	31,235	44.6	43.1	12.3	27.8	27.6	69.1	11.95	0.634	0.714	Stage 3
Kazakhstan	1,916	34.7	38.4	27.0	32.7	25.3	70.1	11.31	0.512	0.305	Stage 3
Kuwait	387	38.5	44.4	17.1	28.4	27.4	65.1	10.91	0.497	0.071	Stage 3
Kyrgyzstan	425	38.4	36.5	25.2	30.8	23.8	74.1	11.11	0.553	0.084	Stage 3
Latvia	606	41.6	48.2	10.2	28.9	25.7	69.8	11.16	0.547	0.136	Stage 3
Lithuania	1,060	43.4	41.6	15.0	24.3	26.8	68.1	11.82	0.637	0.218	Stage 3
Malaysia	4,446	43.4	30.9	25.7	30.9	27.7	68.3	12.11	0.600	0.428	Stage 3
Netherlands	8,188	48.0	41.5	10.5	25.6	29.2	67.1	12.12	0.668	0.518	Stage 3
New Zealand	2,899	46.6	43.7	9.6	26.7	30.9	67.2	11.96	0.676	0.366	Stage 3
Norway	2,374	39.7	43.8	16.4	25.0	25.9	68.1	12.12	0.633	0.336	Stage 3
Oman	363	45.2	29.8	25.1	34.4	27.3	70.8	11.51	0.551	0.061	Stage 3
PR	508	32.5	44.7	22.6	31.7	29.3	69.3	11.38	0.584	0.111	Stage 3
PS	269	43.1	31.2	25.3	27.9	27.9	62.5	12.45	0.575	0.017	Stage 3
Portugal	4,032	46.5	38.3	15.2	27.7	29.4	67.7	11.79	0.635	0.414	Stage 3
Qatar	895	40.8	34.2	25.0	32.6	30.3	67.3	11.78	0.589	0.194	Stage 3
RE	255	38.4	40.4	21.2	26.7	34.1	65.5	11.28	0.660	0.009	Stage 3
Romania	3,252	45.0	39.1	15.9	27.5	24.2	66.8	11.85	0.540	0.383	Stage 3
Singapore	5,094	40.8	33.0	26.2	28.2	27.9	69.5	12.83	0.699	0.448	Stage 3
Slovakia	1,037	42.3	41.4	16.3	27.5	21.7	66.2	11.81	0.486	0.215	Stage 3
Slovenia	673	45.6	41.9	12.5	26.4	24.4	68.5	12.05	0.593	0.152	Stage 3
Spain	14,095	44.9	37.2	17.9	29.2	28.6	66.4	11.83	0.588	0.598	Stage 3
Sweden	4,078	44.1	42.8	13.0	22.9	27.0	68.8	12.08	0.683	0.416	Stage 3
Switzerland	5,097	46.6	41.3	12.1	27.7	28.0	67.6	12.01	0.624	0.448	Stage 3
Taiwan	7,729	45.8	35.1	19.1	27.9	26.6	68.7	12.49	0.646	0.509	Stage 3
Trinidad and Tobago	239	38.5	38.1	23.4	37.2	25.1	69.0	11.40	0.448	0.000	Stage 3
Turkey	8,802	42.2	40.7	17.1	27.3	25.6	65.3	11.70	0.534	0.529	Stage 3
United Kingdom	30,707	44.2	42.0	13.8	27.4	30.5	67.6	12.22	0.684	0.712	Stage 3
United States	219,666	41.9	40.4	17.6	31.9	26.7	68.2	12.31	0.584	1.000	Stage 3
Uruguay	658	41.3	35.7	22.9	27.7	25.4	67.2	11.95	0.570	0.148	Stage 3
Vietnam	10,293	48.0	27.6	24.4	32.5	27.9	66.4	12.21	0.565	0.551	Stage 3